Learning zero-cost portfolio selection with pattern matching

Tim Gebbie[a,b], Fayyaaz Loonat[a,b]

[a]School of Computer Science and Applied Mathematics, University of the Witwatersrand, Johannesburg, South Africa
[b]QuERILab - Quantifying Emergence, Risk and Information

arXiv:1605.04600v1 [q-fin.CP] 15 May 2016

Abstract

We consider and extend the adversarial agent-based learning approach of Györfi *et al* to the situation of zero-cost portfolio selection implemented with a quadratic approximation derived from the mutual fund separation theorems. The algorithm is applied to daily sampled sequential Open-High-Low-Close data and sequential intraday 5-minute bar-data from the Johannesburg Stock Exchange (JSE). Statistical tests of the algorithms are considered. The algorithms are directly compared to standard NYSE test cases from prior literature. The learning algorithm is used to select parameters for agents (or experts) generated by pattern matching past dynamics using a simple nearest-neighbour search algorithm. It is shown that there is a speed advantage associated with using an analytic solution of the mutual fund separation theorems. It is argued that the expected loss in performance does not undermine the potential application to intraday quantitative trading and that when transactions costs and slippage are considered the strategies can still remain profitable when unleveraged. The paper demonstrates that patterns in financial time-series on the JSE can be systematically exploited in collective but that this does not imply predictability of the individual asset time-series themselves.

Keywords: online learning, pattern matching, portfolio control, algorithmic portfolio selection
PACS: 89.65.Gh, 02.50.Ey
2000 MSC: 91-04, 91G10, 91G80
JEL: G11 G14 G17 055

1. Introduction

Sequential investment strategies aim to facilitate portfolio control decisions by collecting information from past behaviour and states of the market and using this information to deploy capital across a selection of assets in a manner the can generate consistent wealth maximization over the long-term [14, 17, 22].

The intention of the paper is not to find a profitable trading strategy for quantitative trading but to show that such strategies exists by providing a simple, transparent and easily recoverable example in the domain of unleveraged zero-cost portfolio selection for statistical arbitrage.

Here we make no specific assumptions relating to the nature of price processes for the sake of the algorithms, however, the approach is broadly based on prior mathematical analysis that use assumptions of stationarity and ergodicity of the price increments in order to allow the study of asymptotic growth rates. In particular to ensure that such growth rates have well-defined maxima when full knowledge of the distribution and its process have been achieved [14, 17, 18, 22].

We investigate the idea that by using pattern matching algorithms (where the patterns are unspecified) combined with learning algorithms, based on some purpose, such as wealth maximisation irrespective or risk [13, 14], we can:

1. Beat a cash portfolio in the context of a self-funding strategy, a zero-cost portfolio strategy, and that
2. We can beat the best stock in the market [16].

The latter has been shown to be the case in prior literature, by investigating daily sampled stock data from the NYSE for long-only (fully invested) portfolio strategies [14, 17, 15, 16, 20]. Here we consider both of these cases: zero-cost, and fully invested strategies, in the context of the South African stock market, the Johannesburg Stock Exchange (JSE), and do so for both daily sampled data and intraday data.

The approach here should not be confused with questioning the value of technical analysis where pre-specified patterns, in the form of some sort of library or set of rules, are used to try to generate systematic wealth [11]. We are considering the problem of probing phenomenology aimed at understanding financial markets as a complex adaptive system [1, 3]. More specifically, we are considering the modelling of time-series arising from complex adaptive systems, something more closely aligned with the context of nonlinear dynamical systems thinking [4]. The question of finding evidence of structure, as opposed to randomness, in financial time-series data, but beyond evidence of long-term memory or typical stylised facts [2]. We argue that we are not trying to show that specific patterns exist, that such pattens are predictable, but rather that the

Email address: tim.gebbie@wits.ac.za (Tim Gebbie)

interaction of a purposeful agent with a stock market using pattern-matching can generate wealth that would not be expected from a typical null-hypothesis of geometric Brownian motion.

We are specifically not looking for statistically preserved properties of time-series, in the sense of time-series models, but are rather looking for evidence of statistical repeating structures in time-series, but without *a-priori* ability to know the form that the structure will take, perhaps because of the nonlinearity of the system in question [8, 10].

We are seeking indirect evidence of structure by showing that a purposeful agent can learn to make investment decisions [3], in a positivist manner, by looking for *a-priori* unspecified and unknown patterns in the data, that can be purposefully exploited, sequentially and systematically, to generate wealth in excess of that expected by randomness and the related normative perspectives of the functioning of financial markets. This is not in itself new, there is a rich literature on attempts at probing the predictability of this or that financial time-series. What can be considered controversial is the view that fairly naive computational learning agents can generate wealth within the system without special insights or understanding of the system itself[1].

By extracting positive growth rates in the excess of the performance of the best stock by using unleveraged combinations of underlying stocks over long periods of time this can be taken as building the case that there are indeed patterns, or some sort of structure, that almost repeat though time in a manner that their occurrence can be treated as exploitable information in collective. This has been shown to be the case for long-only portfolio's [14, 17, 15, 16, 20]. We show this for self-funding strategies; zero-cost portfolio's.

To achieve this we construct sequential investment strategies based on pattern matching and demonstrate that these strategies can generate positive growth rates in excess of the best stocks in the investment universe, and substantial positive growth rates for zero-cost strategies in excess of that expected from investment in cash or risk-free assets.

We do not address the question of whether it is risk that the investor is being compensated for, or even whether the strategies we are isolating are in fact statistical arbitrages, in the sense that the strategies long-term volatility tending to zero in conjunction with an always positive probability of positive performance at zero initial cost [12].

The appearance of patterns and organisation is a fundamental property of complex adaptive systems [4]. Looking directly for pockets of predictability in complex dynamical systems [5] as an approximation to modelling complexity adaptive systems [4] is notoriously difficult given the intricacies of noise and nonlinearity [6, 7]. Coupling purpose, via a learning criterion, here wealth maximisation irrespective of risk, to the selection for patterns, in order to achieve the stated purpose, is the approach promoted here.

It is in this sense that we built a framework that extracts pockets of predictability, if they exist, via pattern searching, ideally in an online manner, in order to increase our agents wealth irrespective of risk, but specifically in the situation where the form of the patterns are always unknown, changing and dynamic, but are represented in collective past histories of the system components.

In Section 2 we present the agent-based learning algorithm as an extension of prior work [13, 14, 17, 18] and [20, 22]. The contributions here are: (i) the algorithm is explicitly re-written in online form in order to make near-real-time applications tractable, (ii) the algorithms are modified for application to the zero-cost portfolio selection problem using the mutual fund separation theorems [25, 24], (iii) the algorithms are explicitly tested, using synthetic data, real daily data both from the NYSE and JSE, and for JSE intraday 5-minute bar-data.

Section 3 describes the approach we have adopted for the generation of experts or agents modified for use in zero-cost portfolio strategies. The algorithm parameters are not tuned prior to use but are left to the online-learning algorithm to select.

In Section 3.6, we consider strategies that target predictable patterns using a simple modified version of the nearest-neighbour pattern-matching strategy developed by [22].

As in the case of the learning algorithm, the agent-generation algorithms have been modified in principle: (i) to support offline and online algorithm use, (ii) they are explicitly framed for use with zero-cost portfolio selection problems, and (iii) portfolio optimizations have been replaced with analytic quadratic approximations in order to improve execution times.

In order to have true online pattern matching the algorithms would have to replaced with either look-up-tables built off-line or a hybrid method that combines offline building of the history of the agents performance and then an almost online method that updates that cached history of agents performance across parameters as the data arrives sequentially in real-time.

Section 4 provides an overview of the data used in the various numerical experiments.

The data is sequential and uniformly sampled and takes on the form of open-high-low-close (OHLC) data, this is described in Section 4.1. The use of open, high, low and close data combinations for the daily data testing can be carried over for intraday studies, and the use of close prices is a special case.

The synthetic data is described in Section 4.2 along with the algorithm testing strategy. Briefly, a simple Kolmogorov-Smirnov test is adopted to assess algorithm behaviour across 4 test cases:

1. SDC1: log-normal random data with zero-means, where no learning should be possible,

[1] This view benefited from conversations with D Hendricks and D Wilcox

2. SDC2: log-normal random data where all assets have the same positive mean and as such basic learning is not possible for zero-cost portfolios (portfolios that have long and short positions that sum to zero),

3. SDC3: log-normal random data with varying positives means, and

4. SDC4: where we have log-normal data with both positive and negative means with the same fixed variance.

The synthetic data is used to understand and prove the behaviour of the zero-cost portfolio strategy (which we will call active portfolio's) and the fully-invested portfolio strategy (which we will call absolute portfolio's).

The four real-world data sets are described in Section 4.3:

1. The standard daily sampled test-data set for the NYSE [14, 17, 15, 16, 20],

2. A more extensive, merged, daily sampled test-data set for the NYSE [32],

3. A daily sampled test-data set for the JSE, and

4. An intraday test-data set for the JSE.

A general overview of the implementation of the numerical experiments is addressed in Section 5.

Section 6 describes the results and analysis of the results, first the synthetic data in section 6.1 and then for the real-world data, in Sections 6.2, 6.3, 6.4, and 6.5, respectively for the four real world case studies: NYSE, extended merge NYSE, daily sampled JSE and intraday JSE.

2. An online-learning algorithm for portfolio selection

The application is for a set of stocks ordered in time where each agent will consider different combinations of stocks for each time-period based on features and strategy parameters. These different agents compete in an adversarial manner in competition for capital allocations [13, 17, 18, 22]. Here agents with poor performance will have incremental capital allocations reduced and agents with robust performance will have incremental increases in capital allocation. Better performing agents will over time have their relative contribution to the aggregate portfolio increased so that their decisions are preferentially selected for trade at the onset of each trading or investment period based on information available at the end of the prior trading period.

The online learning algorithm takes as inputs: a set of agents controls, and performances. These are enumerated over features (here price-relatives) and free-parameters of the temporally ordered objects (here stocks).

The key feature used will be price relatives which are defined for the m-th object as:

$$x_{m,t} = \frac{p_{m,t}}{p_{m,t-1}} \tag{1}$$

In vector notation we will write this equivalently as $\boldsymbol{x}_t$ where the m-th component is $x_{m,t}$.

The controls that represent the agents' are the portfolio weights by which each agent's decision will contribute to the final aggregate decision at a particular time.

Agent performance is represented by factor (agent) mimicking portfolios that are formed from the portfolio controls at each time period. The controls are estimated and implemented at the beginning of each period. The relative changes in asset performance will then modify the relative weights of the asset over the investment period and the performance of a given agent is then determined at the end of the investment period.

This is determined both by the controls, and selecting for the collection of objects the agent is holding, their weights, and the performance of those objects as determined by price relatives.

Agents do not have to hold the same number of objects. Agents can hold all or small groups of objects, they can short-sell objects and hold long positions in objects[2]. The collection of objects a particular agent holds will be called the agent's object cluster.

The parameters that denote agents are typically a parameter that is an index of the cluster of objects an agent has decided to use, and the algorithm specific parameters; typically a data window parameter k determining how much past data to include, and a parameter more specific to a given algorithm if it is required, such as a partition parameter ℓ, and a forecast horizon dependent parameter, τ.

Any four useful parameters can be used in the learning algorithm that was implemented in this paper. The number of agents is then a function of these four free-parameters. The learning algorithm will then carry out the weighted averaging process based on agent past performance over the agents enumerated by these four parameters.

The parameters are denoted τ, w, k and ℓ respectively. We reserved parameters k and ℓ for algorithm specific parameters - this is done in order to try to align with their usage in the prior literature [20]. There are at most W values of w, K values of k, L values of ℓ and τ_n values for the horizon parameter τ.

The default value of the horizon parameter is 1: $\tau = 1$. For simplicity and computational speed the results pre-

[2]Short-selling is when an asset is borrowed for a small fee, and the capital raised from the sale can then be used for other trading or investment activities, for example, the raised capital can be used to buy another asset by taking a long-position. The combination of long and short positions can be cash-neutral where the total value of the initial portfolio is zero. Such a portfolio is called a zero-cost or cash-neutral portfolio.

sented in this paper have used the default value [3]. The choice of these parameters will determine the number of agents in the system. The number of agents is denoted by n where the total number of agents will then be $N = \tau_n WKL$.

The n-th agent is represented by a tuple containing the controls at a given time and its performance $(H_{nm,t}, S_{n,t})$. This tuple will usually be represented in vector notation as $(\boldsymbol{H}_{n,t}, S_{n,t})$ where the object index m is suppressed.

For discrete values of sequential time running from $t = 1$ until some maximal time T the agent controls $\boldsymbol{H}$ are then collection of T time-ordered (N, M)-dimensional matrices that are represented as multi-dimensional double precision matrices in the software.

The value of the n-th agents controls for the m-th object at time t is $H_{nm,t}$ for discrete values of time. The performance of the agents is represented as a (N, T)-dimensional matrix where the n-th agent has its performance over the t-th time interval as $S_{n,t}$.

There are at most M objects. So m can take on values on the integer interval $[1, M]$ that would enumerate the objects. The number of objects remain static for a given agent even though they may be able to achieve zero positions in a particular agent.

From the perspective of the learning algorithm the mechanism of agent generation is not important, it is required that all N agents are correctly enumerated at each time increment. At the beginning of each time increment the controls determined at the end of the previous time increment are implemented and then held to the end of the time period at which time the agent performance is determined and the agent controls are then adjusted using the learning algorithm.

The learning algorithm updates the agent mixture control $q_{n,t}$ which is a measure of how much a given agent will contribute to the aggregate portfolio. The q variables control the relative mixture of agents through time as they compete based on their past performance. The mixture controls cannot in general be thought of as probabilities, which makes their use and notation different to some of the prior literature [20].

2.1. Online-learning algorithm

The learning algorithm is inspired by the universal portfolio approach developed by [17, 18] and refined by [22]. The learning agent can be thought of as a multi-manager, using asset management language, where the multi-manager is selecting and aggregating underlying strategies from a collection of portfolios $\boldsymbol{H}_{n,t}$ and then aggregating using some selection method to a single portfolio $\boldsymbol{b}_t$ that is implemented at each investment or trading period t.

[3] It is anecdotally noted that there is an advantage in learning for the horizon parameter but this does not change the basic point made in this paper

The basic learning algorithm was incrementally implemented online, but offline it can be easily parallelized across agents. The learning algorithm has five key steps:

1. **Update the portfolio wealth**: The portfolio controls $b_{m,t}$ for the m-th asset are used to update the portfolio returns for the t-th period

$$\Delta S_t = \left[\sum_m b_{m,t}(x_{m,t} - 1)\right] + 1 \quad (2)$$

$$S_t = S_{t-1}\Delta S_t. \quad (3)$$

Here the price relatives for the t-th period and m-th asset, $x_{m,t}$, are combined with the portfolio controls for the period just ending to compute the realised portfolio returns for this period, period t. The portfolio controls were computed at the end of the prior period and implemented at the beginning of the current period. The relative amounts of each object in the portfolio will have changed by the relative price changes assuming no cash-flows into or out of the portfolio during this investment period.

2. **Update agent wealth**: The agent controls $H_{nm,t}$ were determined at the end of time-period $t - 1$ for time period t by some agent generating algorithm for N agents and M objects about which the agents make expert capital allocation decisions. At the end of the t-th time period the performance of each agent, $S_{n,t}$, can be computed from the change in the price relatives $x_{m,t}$ for the each of the M objects in the investment universe considered using the prices at the start, $p_{m,t-1}$, and the end of the t-th time increment, $p_{m,t}$, using the agent controls.

$$\Delta S_{n,t} = \left[\sum_m H_{nm,t}(x_{m,t} - 1)\right] + 1. \quad (4)$$

$$S_{n,t} = S_{n,t-1}\Delta S_{n,t}. \quad (5)$$

3. **Update agent mixtures**: We considered three different agent mixture update rules: 1.) the universally consistent choice, and 2.) an exponential gradient choice [23] and 3.) an exponentially weighted moving average. We generically refer to these online updates as rule g. In practice one would select one of the three update rules once for the duration of the offline training, if one seeks to initialise the algorithm prior to deployment, or for use online during the system implementation in real-time. For the numerical experiments presented here we adopted the universal consistent approach inspired by [18, 22] as this demonstrates the principle. We can define the mixture of controls as the accumulated agent wealth is used as the update feature for the next unrealised increment with some normalisation, as such, the agent mixture

control for the n-th agent for the next time increment, $t+1$, is proportional to the measure of wealth:

$$q_{n,t+1} = S_{n,t}. \tag{6}$$

the alternative choices can include the Exponential Gradient (EG)[4] approach of [23] or an Exponential Weighted Moving Average (EWMA)[5] based learning strategy. We adopt the simplest update rule for the mixture of controls, it should be noted that there can be practical advantages to using more adaptive methods such as EG and EWMA learning where the learning rates can be used as additional parameters to be learnt using a thick modelling framework [15].

4. **Re-normalise agent mixtures**: If the agent mixture is to be considered a positive probability then we require that $\sum_n q_n = 1$ and that all $q_n \geq 0$. This is the case of fully-invested agents where no shorting is allowed. We will call these types of agents *absolute* agents:

$$q_{n,t+1} = \frac{q_{n,t+1}}{\sum_n q_{n,t+1}}. \tag{7}$$

For agents that we will consider *active* the leverage is set to unity for zero-cost portfolios: (1.) $\sum_n q_n = 0$ and (2.) $\nu = \sum_n |q_n| = 1$. Here the mixture controls allow for shorting of one agent against another and the portfolio becomes self-funding. The mixture controls can no-longer be thought of as positive probabilities.

$$q_{n,t+1} = \frac{q_{n,t+1} - \frac{1}{N}\sum_n q_{n,t+1}}{\sum_n |q_{n,t+1} - \frac{1}{N}\sum_n q_{n,t+1}|} \tag{8}$$

The leverage is normalised in order to ensure consistency between the learning algorithms and agent generating algorithms.

5. **Update portfolio controls**: The portfolio controls $b_{m,t}$ are updated at the end of time period t for time period $t+1$ using the agent mixture controls $q_{n,t+1}$ from the updated learning algorithm and the agent controls $H_{nm,t+1}$ from the agent generating algorithms using information from time period t and averaged over all n agents.

$$b_{m,t+1} = \sum_n q_{n,t+1} H_{nm,t+1}. \tag{9}$$

[4] *Exponential Gradient* (EG) based learning:

$$q_{n,t+1} = q_{n,t} e^{\left(\frac{\eta S_{n,t}}{\sum_n q_{n,t} S_{n,t}}\right)}$$

[5] *Exponential Weighted Moving Average* (EWMA) based learning:

$$q_{n,t+1} = \lambda q_{n,t} + (1-\lambda)\left(\frac{q_{n,t} S_{n,t}}{\sum_n q_{n,t} S_{n,t}}\right)$$

The strategy is to implement the portfolio controls, wait until the end of the increment, measure the features, update the agents and then re-apply the learning algorithm to compute the agent mixtures and portfolio controls for the next time increment.

Algorithm 1 Online-Learning Algorithm (OLA)

Require:

1. updated agent-controls $\boldsymbol{H}_{n,t+1}$
2. current feature realisation $\boldsymbol{x}_t$
3. current portfolio controls $\boldsymbol{b}_t$
4. current agent-controls $\boldsymbol{H}_{n,t}$
5. past agent-wealth $S_{n,t-1}$
6. past portfolio wealth S_{t-1}

for t-state **do**
Step 1: The portfolio wealth is updated
$S_t = S_{t-1}(\boldsymbol{b}_t(\boldsymbol{x}_t^T - 1) + 1)$
Step 2: The agent wealth is updated
$S_{n,t} = S_{n,t-1}(\boldsymbol{H}_{n,t}(\boldsymbol{x}_t^T - 1) + 1)$
Step 3: The agent mixture is updated for rule g
$q_{n,t+1} = g(q_{n,t}, S_{n,t})$
Step 4: The agent mixtures are re-normalised
$q_{n,t+1} = \begin{cases} \sum_n q_{n,t+1} = 1,\ q_{n,t+1} \geq 0 \\ \sum_n |q_{n,t+1}| = 1, \sum_n q_{n,t+1} = 0. \end{cases}$
Step 5: The portfolio controls are updated
$\boldsymbol{b}_{t+1} = \sum_n q_{n,t+1} \boldsymbol{H}_{n,t+1}$
Leverage corrections
if $(\nu = \sum_m |b_{m,t}|) \neq 1$ **then**
renormalise controls
$\boldsymbol{b}_{n,t+1} = \frac{1}{\nu}\boldsymbol{b}_{n,t+1}$
renormalise mixtures
$q_{n,t+1} = \frac{1}{\nu} q_{n,t+1}$
end if
end for
return ($\boldsymbol{b}_{t+1}$,$S_{n,t}$,S_t,$q_{n,t+1}$)

3. Agent generating algorithms

The purpose of the agent generating algorithms are to sequentially generate the agent controls $H_{nm,t}$ for the n-th agent for the m-th object for implementation at the start of the t-th time period. These will be denoted in vector notation as $\boldsymbol{H}_{n,t}$.

We initially considered three different agent-generating algorithms over which the thick modelling was carried out in order to learn the various algorithms' free-parameters: 1.) a pattern-matching algorithm [20], 2.) a contrarian mean-variance portfolio algorithm we called anti-BCRP (as it trades against the Best Constant Rebalanced Portfolio for a given k-tuple of data) [6], and 3.) the ANTICOR

[6] The anti-BCRP algorithm can be used to learn for mean-

algorithm [16]. The various free-parameters of these algorithms, such as the window sizes k and partitions ℓ were then used to enumerate the agents that would compete for capital allocations in the learning algorithm.

We adopted the pattern-matching approach [20] for the numerical experiments in this paper as we found a performance advantage in looking for more general patterns rather than merely targeting mean-reversion effects, and more importantly, the pattern-matching algorithms are more generic as they do not require any *a-priori* choices for the structures that are learnt for. This was considered to be more faithful to the intent of the paper - where we seek to show that unspecified patterns can be learnt for in a manner that can both beat the best single stock in a universe of stocks and can beat a cash portfolio in a self-funding strategy.

3.1. Comments on Notation

The feature realisations at time t for the m-th object, $x_{m,t}$, are also denoted in vector notation as $\boldsymbol{x}_t$. The agent controls and the feature time-series are the key inputs in the online-learning algorithm to determine the agent mixtures $q_{n,t}$ through time. The online learning algorithm is path-dependent and as such both a function of the history of agent controls as well as the feature time-series history.

Following prior work we denote random feature variables as $\boldsymbol{X}$ and their realisations as $\boldsymbol{x}$ [20, 22] where for some vector valued stationary and ergodic process $\{\boldsymbol{X}_t\}_{-\infty}^{+\infty}$ with realisations denoted as $\boldsymbol{x}_1, \boldsymbol{x}_2, \ldots, \boldsymbol{x}_t$ and their corresponding random variables as $\boldsymbol{X}_1, \boldsymbol{X}_2, \ldots, \boldsymbol{X}_t$. However, we will refine the notation further in order to more effectively enumerate the agents for our specific implementation.

The strategies are based on constructing a k-tuple of the selected feature for m-objects. We will denote the agent-tuple by $\boldsymbol{x}_{k\ell w,t}$ and the k-tuple as $\boldsymbol{x}_t^{t-k}$. The k-tuple is a slice of data of length k from the current time t, of width m enumerating all the objects. We will modify the k-tuple notation to $\{\boldsymbol{x}_t^{t-k}\}_{s(n),\ell}$ to denote a k-tuple taken from an ℓ-partition of the data for a given cluster of objects $w = s(n)$. Here s is the cluster index of the n-th agent. We are suppressing the m index and using vector notation to write the k-tuple as $\boldsymbol{x}$. The agent-tuple will be unique to the n-th agent where n is the unique agent index enumerating a particular combination of k,ℓ and w.

A k-tuple is used to determine agent controls $\boldsymbol{H}_{n,t}$. The initial features used are historical prices sequences which are assumed to be realisation $\boldsymbol{x}$ from some random process $\boldsymbol{X}$. The pattern-matching algorithm will then refine the k-tuple to groups of nearest-neighbours that are expected to reflect historical selected outcomes that better reflect future outcomes than merely the last price change or price change sequence. This is done by comparing the current realisation $\boldsymbol{x}_t^{t-k}$ with the past.

In this way, given a set of parameters enumerating the n-th agent we will select the required tuple from the existing data realisations depending on the algorithm parameters using some selection function f

$$\boldsymbol{x}_{n,t} = \boldsymbol{x}_{n(k,\ell,w),t} = \boldsymbol{x}_{k\ell w,t} = f_{\ell,w}(\boldsymbol{x}_1^t, \boldsymbol{x}_t^{t-k}) \tag{10}$$

where the m-th component of the k-tuple is $x_{n,t}^m$.

3.2. The Log-optimal strategy

The log-optimal strategy under the assumptions of stationarity and ergodicity has been shown to be the best possible choice of strategy over the long term [14]. This type of analysis has been extended to the semi-log-optimal case [20] where weakened conditions have been derived.

The surprising result is that even with this weaker formulation the loss of optimality is such that log-optimality has, for all practical purposes, equivalent performance to portfolios selected using semi-log-optimality [20]. This provides an argument for the use of competing sequences of mean-variance portfolios in the framework of agent-based competition for capital.

With an initial investment wealth of S_0 using a sequence of portfolio controls $\boldsymbol{B} = \{\boldsymbol{b}_i\}_{i=1}^{t-1}$ from time $i = 1$ until the current time t the portfolio wealth for a fully-invested portfolio is [20]

$$S_t = S_0 \Pi_{i=1}^T \boldsymbol{b}(\boldsymbol{x}_1^{i-1})\boldsymbol{x}_i^T = S_0 e^{\sum_{i=1}^T \log(\boldsymbol{b}(\boldsymbol{x}_1^{i-1})\boldsymbol{x}_i^T)}. \tag{11}$$

This gives an average portfolio growth rate $W_t(\boldsymbol{B}) = \frac{1}{T}\sum_{i=1}^T \log(\boldsymbol{b}(\boldsymbol{x}_1^{i-1})\boldsymbol{x}_i^T)$. The log-optimal portfolio selection problem is thus

$$\boldsymbol{b}^*(X_1^{t-1}) = \arg\max_{\boldsymbol{b}} \mathrm{E}\left[\log(\boldsymbol{b}(X_1^{t-1})\boldsymbol{X}_t)|\boldsymbol{X}_1^{t-1}\right]. \tag{12}$$

Here one is aiming to maximize the overall wealth through the incremental selection of the sequence of fully-invested portfolio controls $\boldsymbol{B}$.

3.3. Universally consistent strategies

The fundamental result of universal log-optimality is that no investment strategy can have a faster average rate of growth than that arising from the log-optimal portfolio [13, 14, 17, 18]. However, full knowledge of the distribution of the process is required. Strategies achieving an equivalent growth rate without knowing the distribution are called *universally consistent* [14, 20] strategies.

In principle one could via simulation enumerate all the possible controls and find via brute-force the set of controls that solve the log-optimal portfolio selection problem.

This is ambitious given current technology constraints and that the opportunity set of stocks is typically large and the data representing the features even larger - particularly for intraday quantitative trading problems.

reversion by directly using k past realisations of performance of each object, for a given partition, by finding the mean-variance wealth minimizing portfolio (in order to be contrarian), either fully-invested or zero-cost, and using the resulting portfolio weights for the agents with the specific window and partition parameters: $\boldsymbol{H}_{n,t+1} = \boldsymbol{H}_{n,t+1}(\gamma, -\boldsymbol{\mu}(\boldsymbol{x}_{n,t}), \Sigma(\boldsymbol{x}_{n,t}))$ comparing with Eqn. (25) and (26).

In the idealized situation we would define some simplex Λ where there is a prior distribution μ on the simplex, such that some expert or agent $\boldsymbol{b}$ is a given realisation from this distribution of portfolios. We would then directly evaluate the μ-weighted fully-invested universal portfolio at time t [18, 19]

$$\boldsymbol{b}_t^* = \frac{\int_\Lambda \boldsymbol{b} S_{t-1}(\boldsymbol{b}, \boldsymbol{x}_{t-1}) d\mu(\boldsymbol{b})}{\int_\Lambda S_{t-1}(\boldsymbol{b}, \boldsymbol{x}_{t-1}) d\mu(\boldsymbol{b})} \tag{13}$$

where $\int_\Lambda d\mu(\boldsymbol{b}) = 1$ and the portfolio value S_t at time t is as

$$S_t(\boldsymbol{b}, \boldsymbol{x}_t) = \prod_{j=1}^{t} \boldsymbol{b}\boldsymbol{x}_j^T = \prod_{i=1}^{t} \sum_{j=1}^{m} b_j x_{j,t}. \tag{14}$$

Here the portfolio is fully-invested such that $\boldsymbol{b}\mathbf{1}^T = 1$ for unit vector $\mathbf{1}$.

Although we seek strategies that are universally consistent with respect to the class of stationary and ergodic processes. A pragmatic approach is required given both the unrealistic distributional assumptions, and the curse of dimensionality we face in enumerating control space [7].

The strategy is to reduce the problem by finding a more informed subset of controls that can be used to approximate the required sequence of portfolio controls that are used to represent a universally consistent strategy. In addition to reducing the set of applicable controls one also aims to streamline the evaluation of these controls and their adaption through time, this can be achieved by reducing the log-optimality criterion to semi-log-optimality.

3.4. Semi-log optimality

We choose to focus on the first two moments of the price relative distributions: the mean and covariance. This will allow enhanced performance speed of the algorithms (see Figure 13) but with some loss in long-term optimality [30, 20] and as such a deviation from the universally consistent strategies.

First, we have reduced the opportunity space in the simplex of all possible portfolios in order to make the problem of finding a portfolio that is optimal over the entire feature space computationally tractable, this is achieve by using agent-generating algorithms and learning over the free-parameters for those agents generating algorithms.

Second, we replaced the optimization with a quadratic approximation that will give us analytic solutions to replace optimizations that we would otherwise have to solve numerically. In addition to a performance advantage, using the quadratic approximation this will also provide a straight-forward method for considering both fully-invested and zero-cost portfolio's in a single framework.

Streamlining the algorithms for performance was approached in two steps, first, to separate the problem into that of an online-learning algorithm and the agent generating algorithms, then, second, to reduce the log-optimality criterion to semi-log-optimality.

The semi-log-optimal portfolio selection takes on the form

$$\boldsymbol{b}^*(X_1^{t-1}) = \arg\max_{\boldsymbol{b}} \mathrm{E}\left[h(\boldsymbol{b}(X_1^{t-1})\boldsymbol{X}_t)|\boldsymbol{X}_1^{t-1}\right]. \tag{15}$$

where $h(z) = (z-1) - \frac{1}{2}(z-1)^2$ from the second order Taylor expansion of $\log(z)$ at $z = 1$.

A related approach was taken in [19] where they derived an analytic approximation for an efficient universal portfolio. Our simplified mean-variance approach was motivated by their development of an analytic algorithm, the difference here is that we want an algorithm that is online, analytic, explicitly includes zero-cost portfolios, and allows for the restriction of the solution space using some agent generating algorithm directly at each step rather than via side-information.

3.5. Active fund separation problem

The determination of the optimal portfolio is sequentially implemented using the exact solution to the quadratic approximation to log-optimality by solving the active fund selection problem. The active fund selection problem is a special case of the mutual fund selection problem [25, 24]. This will give an analytic approximation that can both cater for long-only fully-invested agents (absolute agents) as well as leverage one[8] zero-cost portfolio's (active agents).

We therefore consider the semi-log-optimal portfolio optimization problem [25, 26, 24] for return expectation vector $\boldsymbol{\mu}$ and asset return covariance matrix Σ with a portfolio control vector $\boldsymbol{\omega}$ in terms of the risk aversion parameter γ. The conjugate transpose of a vector is denote as $(\cdot)^T$ over a single investment period to define the control problem as:

$$\max_{\boldsymbol{\omega}} \left\{ \boldsymbol{\omega}^T \boldsymbol{\mu} - \frac{\gamma}{2} \boldsymbol{\omega}^T \Sigma \boldsymbol{\omega} \right\} \text{ s.t. } \boldsymbol{\omega}^T \mathbf{1} = 1. \tag{16}$$

Here we have changed notation to denote the portfolio controls as $\boldsymbol{\omega}$ in order to avoid confusion with the portfolio strategy controls $\boldsymbol{b}$ that are the result of the online-learning algorithm which aims to approximate the semi-log-optimal portfolio selection strategy for aggregate portfolio controls $\boldsymbol{b}_t$ for time increment t.

Here the portfolio controls $\boldsymbol{\omega}$ are used to generate the agents that populate the agent control set $\boldsymbol{H}_{n,t}$. It is the agent control set that is then used to generate the semi-log-optimal portfolio choice at each time t: $\boldsymbol{b}_t$.

Eqn. (16) can be rewritten as the mutual-fund Lagrangian

$$L = \boldsymbol{\omega}^T \boldsymbol{\mu} - \frac{\gamma}{2} \boldsymbol{\omega}^T \Sigma \boldsymbol{\omega} - \lambda_\omega (\boldsymbol{\omega}^T \mathbf{1} - 1). \tag{17}$$

[7] For each random process in the long-term limit the growth-rate of these strategies is equivalent to that of the log-optimal portfolio when full-knowledge of the distribution is available. In order to construct such universally equivalent strategies one needs to know the conditional distribution $\boldsymbol{X}_t$ given some past $\boldsymbol{X}_1^{t-1}$.

[8] $\sum_i |\omega_i| = 1$ for portfolio controls ω.

and solved using elementary Kuhn-Tucker methods. Two equations are found in terms of the optimal solution for the portfolio control, $\boldsymbol{\omega}^*$, the first gives the quadratic optimal risk-return pay-off, and the second, the fully-invested portfolio investment constraint

$$\begin{aligned} \boldsymbol{\omega}^* &= \frac{1}{\gamma}\Sigma^{-1}\left(\boldsymbol{\mu} - \lambda_\omega \mathbf{1}\right), & (18)\\ \boldsymbol{\omega}^{*T}\mathbf{1} &= 1. & (19) \end{aligned}$$

The Lagrange multiplier is determined by substituting Eqn. (18) into Eqn. (19) to find:

$$\lambda_\omega = \frac{\mathbf{1}^T\Sigma^{-1}\boldsymbol{\mu}}{\mathbf{1}^T\Sigma^{-1}\mathbf{1}} - \frac{\gamma}{\mathbf{1}\Sigma^{-1}\mathbf{1}}. \quad (20)$$

This is then used to eliminate the Lagrange multiplier from Eqn. (18) to find a formulation of the mutual fund separation theorem:

$$\boldsymbol{\omega}^* = \frac{\Sigma^{-1}\mathbf{1}}{\mathbf{1}^T\Sigma^{-1}\mathbf{1}} + \frac{1}{\gamma}\Sigma^{-1}\left(\boldsymbol{\mu} - \mathbf{1}\frac{\mathbf{1}^T\Sigma^{-1}\boldsymbol{\mu}}{\mathbf{1}^T\Sigma^{-1}\mathbf{1}}\right). \quad (21)$$

The first term on the right is the lowest risk portfolio and the second term is the zero-cost portfolio that encapsulates the relative views of the assets. We will typically work with the separation theorem in the form given in Eqn. (21). The second term will give us an efficient method of generating zero-cost portfolio's.

It is then convenient to re-write the Mutual Fund Separation theorem to an Active Fund Separation theorem explicitly from Eqn. (21) by defining the lowest risk portfolio as the benchmark portfolio:

$$\boldsymbol{\omega}^* = \boldsymbol{\omega}_{_B} + \boldsymbol{\omega}_{_A}, \quad (22)$$

where

$$\begin{aligned} \boldsymbol{\omega}_{_B} &= \frac{\Sigma^{-1}\mathbf{1}}{\mathbf{1}^T\Sigma^{-1}\mathbf{1}}, & (23)\\ \boldsymbol{\omega}_{_A} &= \frac{\Sigma^{-1}}{\gamma}\left(\frac{\boldsymbol{\mu}\mathbf{1}^T - \mathbf{1}\boldsymbol{\mu}^T}{\mathbf{1}_T\Sigma^{-1}\mathbf{1}}\right)\Sigma^{-1}\mathbf{1}, & (24) \end{aligned}$$

The formulae for $\boldsymbol{\omega}_{_B}$ and $\boldsymbol{\omega}_{_A}$ will be directly used in the agent generating algorithms based on views encoded in the mean, $\boldsymbol{\mu}$, and the covariances, Σ, as a function of the various agent generating parameters. The resulting controls $\boldsymbol{H}_{n,t}$ will then be determined from the m-th component of either $\boldsymbol{\omega}_{_A}$ for the active agents or $\boldsymbol{\omega}_{_B} + \boldsymbol{\omega}_{_A}$ for the absolute agents for the n-th agent for time-increment t.

For situations where we want agents constructed from zero-cost portfolios we will use the tactical solution from Eqn. (24) to generate the agents for a given k-tuple. In situations where we need fully invested agents we will use the combination of the benchmark fund and the active (or tactical) fund.

Suppressing indexes over the m objects the agent controls for the n-that agent for the two possible cases: (1.) the absolute agents, and (2.) the active agents is then

$$\boldsymbol{H}_{n,t} = \begin{cases} \boldsymbol{h}^T\mathbf{1} = 1, \boldsymbol{h} = \boldsymbol{\omega}_{_B}(\Sigma) + \boldsymbol{\omega}_{_A}(\gamma, \boldsymbol{\mu}, \Sigma) \text{ s.t. } \boldsymbol{h} \geq 0 \\ \boldsymbol{h}^T\mathbf{1} = 0, \boldsymbol{h} = \boldsymbol{\omega}_{_A}(\gamma, \boldsymbol{\mu}, \Sigma) \text{ s.t. } \boldsymbol{h}^T\boldsymbol{h} = 1. \end{cases} \quad (25)$$

Here the m-th component of $\boldsymbol{H}_{n,t}$ is $H_{nm,t}$ and the portfolio weights are dependent on the agent-tuples $\boldsymbol{x}_{n,t}$ for a given agent

$$\begin{aligned} \boldsymbol{\omega}_{_A} &= \boldsymbol{\omega}_{_A}(\gamma, \boldsymbol{\mu}(\boldsymbol{x}_{n,t}), \Sigma(\boldsymbol{x}_{n,t})) & (26)\\ \boldsymbol{\omega}_{_B} &= \boldsymbol{\omega}_{_B}(\Sigma(\boldsymbol{x}_{n,t})). & (27) \end{aligned}$$

For the active agent we enforce the leverage unity constraint at the beginning of each time increment, this can be considered equivalent to setting the risk-aversion γ, at the beginning of each time increment, such that the leverage is always unity.

This is an important feature of the algorithm as we do not enforce uniform risk-aversion through time. We rather choose to ensure that capital be fully utilized given the available information. The following sections describe how the agent-tuples are constructed for the various agent generating algorithms.

3.6. Agent generating algorithms from patterns

In order to efficiently reduce the space of portfolio controls to efficiently generate a reasonable approximation to universally consistent strategies using Eqn. (13) we reduce the set of applicable controls using agent-generating algorithms. The agent-generating algorithm we use in our numerical experiments will be a pattern-matching algorithm [20]. One can make various decisions about how to break data up into manageable pieces for the various algorithms, the most basic decisions relate to how to break up the data in time, we call this partitioning, the other choice relates to how we break the data up in terms of the objects themselves (often called the features), this we call clustering. Partitioning is typically a more intricate task because this has implications for the algorithm and system structure.

The pattern-matching algorithm is based on two steps subsequent to the choice of clusters $s(n)$: (1.) partitioning and (2.) pattern-matching. Clusters can be chosen by a variety of methods, we would like to promote two methods: (i) correlation matrix based methods [27], and (ii) clusters based on economic classifications of stocks [9]. The prior method, correlation based methods, have outputs that can be directly used as inputs into the algorithms discussed here, specifically via $s(n)$, the cluster membership parameters. It is however, the method based on fixed economic sector classifications [29], that will be explicitly used in this paper for the intraday experiments in Section 6.5, this is both for speed and simplicity [10].

In the daily numerical experiments we have ignored the impact of clustering and used the clusters $s(n)$ of the n-th

[9] For example, using ICB (Industry Classification Benchmark) sectors classifications [29]

[10] It should be noted that using ICB sectors to generate additional agents for the daily simulations does boost algorithm wealth performance but we chose to explicitly demonstrate the value of including sector information in the context of the intraday strategies

stock as being trivial, *i.e.* we consider a single stock cluster that includes all m objects. The inclusion of clustering indexing can be important to the practical implementation of these techniques as it is often useful to restrict trading signal decisions to similar stocks. There is a wealth advantage to this, as we have shown when we considered the impact of clustering for the numerical experiments using intraday data (see Table 15).

The pattern matching algorithm is split into two key components: First, the *partitioning algorithm*, which selects collection of time-order features from the full set of feature data. Second, the *pattern-matching algorithm*, where given a measured pattern derived from the feature data, is used by the algorithm to find similar patterns in a given partition of the feature data.

3.6.1. Partitioning

Subsets of time-ordered data are selected from the original time-order data for a given collection of objects. The collection of objects can in turn be a sub-collection of the original set of objects. Partitioning takes place in the time domain while clustering is in the object dimension. The purpose of partitioning is to prepare data subsets for pattern-matching [22]. Four distinct approaches to data partitioning are enumerated here, however only the trivial partition is used in the experiments.

A partition is a collection $\{p_t\}_\ell$ represented by a logical vector of the length of a given time-series where true is represented as one and false as zero to index membership in a given partition. When a partition is determined from features that determine the state of the system at a given time we will use that partition to represent the system in that state for the sake of pattern-matching.

For the numerical experiments presented here we will use variations of the *trivial partition*: Here all the temporally ordered data is kept in a single partition as represented by a vector of ones of length of the time-series.

$$\{p_t\}_1 = \{(1, \ldots, 1, 1, 1)\}. \tag{28}$$

There are wealth advantages associated with more sophisticated partitions. We considered four different partitioning approaches: the *trivial partition*, the *over-lapping partition* [11]: were data membership in partitions is repeated in order to bias the data towards a given time, for example, the last time-increments is repeated across all ℓ partitions for time-series of length T, the *exclusive partition* where the partitions are mutually exclusive subsets of the full partition, and the *side-information partition* [18].

The most heuristically useful partition is that of the *side-information partition* where partitions can be preselected in the partitioning algorithm based on rules conditioned on side-information [18], partitioning can be both useful as a nuanced exploitation of information, for example by splitting feature data over different regimes, and thus to generate distinct agents for different regimes, and as an effective approach to parallelization of algorithms.

Here we would partition the time-series based on side-information arising from additional features drawn from the system being observed as in [18]. For example, we could use a Markov-switching algorithm with ℓ states, assign each time in the time-series a state index and the define the partition membership based on states, or we could choose a feature as side-information and ℓ-tile the data into ℓ groups and then based on whether a given time has a side-information feature in a particular group it would be assigned to a given partition.

Partitioning serves as a convenient mechanism for breaking up the feature data into distinct states. This can be useful when choosing to search for patterns when the system is in a distinct state as it will enable the algorithm to search for patterns only in historic data residing from times in the past when the system was in a similar state. By combining a partitioning algorithm with a state-detection algorithm one can both improve computational times as well as algorithm performance in terms of wealth generation [28], this is not explored further here.

3.6.2. Pattern-matching

The pattern-matching algorithm will take a k-tuple and search a given partition of the feature data for similar patterns by finding the smallest distance measure between the k-tuple and data in a given partition. This best matching set of data in the partition will then be used to determine a pattern-matching time j_ℓ. The matching time will then be used to select a future outcome some time period τ ahead of the matched pattern. This future outcome is used to construct a tuple of data, the agent-tuple, iteratively using the look-ahead rule: $j_n = j_\ell + \tau$. A number of such pattern-matches will be accumulated to construct the agent-tuple $\boldsymbol{x}_{n,t}$ and from this a mean and covariance are computed.

This mean and covariance will then serve as the input into Eqn. (22) to determine that agent controls $\boldsymbol{H}_{n,t+1}$, the n-th agents controls to be held for time-period $t+1$.

The pattern-matching algorithm is split into two separate algorithms. The first algorithm, which we will call the pattern algorithm, generates patterns to be matched and partitions of data into which the pattern will be matched. The second algorithm will then take the pattern and the data partitions and generate matching times. The matching times will then be used to generate an agent-tuple $\boldsymbol{x}_{n,t}$.

The pattern algorithm generates a k-tuples $\{\boldsymbol{x}^t_{t-k}\}_{s(n)}$ [22] for matching, and a data partition $\{\boldsymbol{x}_t\}_{(p_\ell, s(n))}$ using a predefined temporal partition $\{p_\ell\}$ of the data and the cross-sectional cluster for the n-th agent $s(n)$. This is iteratively done for each agent as enumerated by the parameters that define a given agent: the cluster membership $w = s(n)$ of the n-th agent, the partition variable ℓ,

[11] Example of length T overlapping partition of features:

$$\{p_t\}_T = \{(0, \ldots, 0, 0, 1), (0, \ldots, 0, 1, 1), \ldots, (1, \ldots, 1, 1, 1)\}.$$

the k-tuple variable k and the look-ahead horizon variable τ.

For each set of variables that define the n-th agent the pattern algorithm will then call the matching algorithm.

Algorithm 2 PATTERN-MATCHING Algorithm (PMA)

Require:

1. features $\boldsymbol{x}_t$
2. n-agent parameters k, ℓ, $s(n)$, τ
3. partitions $\{p_\ell\}$

for n-agents **do**
$\boldsymbol{H}_{n,t+1}$=MATCHING$(\tau,\{p_\ell\},\{\boldsymbol{x}^t_{t-k}\}_{s(n)},\{\boldsymbol{x}_t\}_{(p_\ell,s(n))})$
end for agents
return $\boldsymbol{H}_{t+1}$

The matching algorithm will find matches for the k-tuples, $\boldsymbol{x}^t_{t-k}$ in the partitions. If there is a single partition of data, the matching algorithm will find the $\hat{\ell}$ closest matches. We consider two rules for calculating $\hat{\ell}$ and will refer to these as rule P. This rule is introduced in order to easily compare our algorithms with prior literature, more specifically [20, 22]. The difference is related to how the partitions are defined and implemented.

We consider the *trivial rule*: $\hat{\ell} = \ell$ and the rule required to recover the Nearest-Neighbour (NN) algorithm performance described in [22]. The *Györfi et al Nearest Neighbour rule* is where $\hat{\ell}$ is determined by a variable $p_\ell \in (0, 1)$. The choice of p_ℓ used in the experiments is the same as in [22].

$$p_\ell = 0.02 + 0.5\frac{\ell - 1}{L - 1} \tag{29}$$

$$\hat{\ell} = \lfloor p_\ell t \rfloor \tag{30}$$

Where t represents the number of time periods in the history, and the floor is taken to find the smallest partition at the given time. This modification serves primarily to allow us to recover prior results in the literature using the framework we implemented in the software for the numerical experiments.

If there are ℓ partitions of data the algorithm will find the best match in each partition. The matching algorithm will find ℓ best matches and from those best matches extract ℓ matching times j_ℓ associated with the time of each k-tuple match. From the look-ahead rule the matching algorithm will then construct the agent-tuple $\boldsymbol{x}_{n,t}$. The matching algorithm will then compute the agent-control for this given agent-tuple $\boldsymbol{h}_{n,t}$.

The distance between tuples is the 2-norm. Although we could use the distance between two matrices as the general distance in the algorithm, we have chosen to differentiate selecting the most recent vectors of object features and the test-tuple as the vector distance between these two vectors only for the case of $k = 1$, while for $k > 1$ we measure the distance of each object from the same object at a difference time independently from other objects.

This will rather allow us to search for the best fits of objects independently rather than in collective. This is an important refinement, in the original version of the algorithm we followed [22] and used the 2-norm in full generality independent of the window size k we found better performance by independently selecting for patterns using column-wise computed distances.

4. Data Description

4.1. OHLC data

The data we will consider will be sequential data, but not necessarily continuously sequential. For this reason we will study OHLC (Open-High-Low-Close) bar-data where the closing price of a given bar is not necessarily the opening price of the subsequent bar of the data. We will first study daily sampled data and then intraday data. The algorithms will be initially tested using synthetic data (see Section 4.2), and then the real world test data used in prior research [17, 22] (See Section 4.3) which are sequences of daily sampled closing prices.

The data and algorithms can be easily extended to accommodate additional features as side-information [18]; such as volumes, spreads, and various financial indicators and asset specific and state attributes. The side-information can be trivially used to re-partition data into additional sets of agents and then used as inputs into the learning algorithm. The wealth performance enhancement relating to the side-information extension is not demonstrated in the numerical experiments presented here.

OHLC bar data is typically represented by a candle-stick graph as in Figure (1).

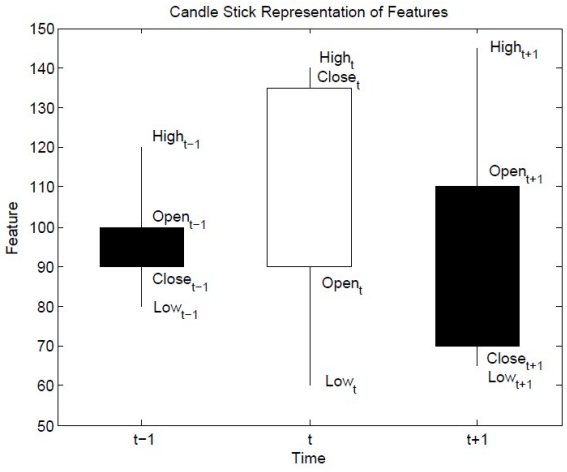

Figure 1: The feature time-series data is best thought of as OHLC (Open-High-Low-Close) bar data. The filled box in the candle chart denotes the situation where the close price is lower than the open price, conversely the unfilled box has the close price higher than the open price.

The time-series data is such that the closing price of time-increment t is not necessarily at time $t + 1$ the start

Algorithm 3 MATCHING Algorithm (MTA)

Require:

1. look-ahead-rule τ
2. partition $\{p_\ell\}$
3. k-tuple $\{\boldsymbol{x}_{t-k}^{t}\}_{s(n)}$
4. data partition $\{\boldsymbol{x}_t\}_{(p_\ell,s(n))}$

for t-state **do**
  for $p_\ell \in \{p_\ell\}$ **do**
    for j-states $\in p_\ell$ **do**
      find a test-tuple
      $\boldsymbol{s}_{t(j)-k}^{t(j)} = \{\boldsymbol{x}_{t(j)-k}^{t(j)}\}_{(p_\ell,s(n))}$
      find distance to k-tuple
      $\boldsymbol{\epsilon}_{k,j} = \{\boldsymbol{x}_{t-k}^{k}\}_{s(n)} - \boldsymbol{s}_{t(j)-k}^{t(j)}$
      if k=1 **then**
        compute the 2-norm for vector $\boldsymbol{\epsilon}_{1,j}$
        $\epsilon_j = \sum_{m\in\text{objects}} \sqrt{\epsilon_{m1,j}2} = \sqrt{\boldsymbol{\epsilon}_{1,j}\boldsymbol{\epsilon}_{1,j}^{T}}$
        distance measure of dim(objects)
        $\{\boldsymbol{\epsilon}_j\}_{p_\ell} \leftarrow \epsilon_{m,j} = \epsilon_j \forall m$
      else
        column-wise 2-norms for matrix $\boldsymbol{\epsilon}_{j,k}$
        $\{\boldsymbol{\epsilon}_j\}_{p_\ell} \leftarrow \epsilon_{m,j} = \sum_{k'=1}^{k} \sqrt{\epsilon_{mk',j}^2}$
      end if
    end for states
    if $\dim(\{p_\ell\}) = 1$ **then**
      Switch NN algorithm partition choice [22]
      $\hat{\ell} = P(\ell, t)$
      Find $\hat{\ell}$ matching-times in a single partition
      $\boldsymbol{j}_\ell = \min_{j\forall \dim(j)=\ell} \{\boldsymbol{\epsilon}_j\}$
    else
      Find the best match in each of the ℓ partitions
      $\boldsymbol{j}_\ell = \min_{j\forall p_\ell\in\{p_\ell\}} \{\boldsymbol{\epsilon}_j\}_{p_\ell}$
    end if
  end for partitions
  update the look-ahead-rule
  $j_n = j_\ell + \tau$
  Update the agent-tuple
  $\boldsymbol{x}_{n,t} = \{\boldsymbol{x}_1^t\}_{t\in j_n}$
  Update the mean and covariance
  $\boldsymbol{\mu} = \boldsymbol{\mu}(\boldsymbol{x}_{n,t} - 1)$
  $\Sigma = \Sigma(\boldsymbol{x}_{n,t} - 1)$
  Update the agent-controls
  $\boldsymbol{H}_{n,t+1} = \boldsymbol{H}_{n,t+1}(\gamma, \boldsymbol{\mu}, \Sigma)$
end for t-state
return $\boldsymbol{H}_{n,t+1}$

of time increment $t+1$. The closing price can in fact be at some time $t+\delta$ for some arbitrary data-specific time-increment δ.

A low-frequency example is that of a typical trading day on the JSE, the market opens in the morning with some opening price, o_t, at 9h00, the market may then close at some closing time 17h00, after a closing auction period, the official closing price c_t, is then printed soon after the market close (perhaps after some randomisation period). The market is then closed for some time-period over-night until the market opens again on the subsequent day. There is a period, δ, when the market is close and as such information is not continuously being priced into the traded assets. Information that accumulates over-night will then be priced into the market prices through the process of the opening auction and subsequent trading in the various assets.

Our approach to OHLC data is applicable to a variety of synchronously sampled or re-sampled data sets, including intraday data:

1. *close-to-close*: Here the prices $p_{m,t}$ for the m-th assets are the time-series of close prices. The price relatives $x_{m,t}$ are then the computed from the close price time-series $c_{m,t}$
$$x_{m,t} = \frac{c_{m,t}}{c_{m,t-1}}. \tag{31}$$
The algorithm is trying to exploit information relating to price changes from the close of trading of one time increment to the close of trading of a subsequent time increment.

2. *open-to-close*: Here the prices $p_{m,t}$ for the m-th assets are the ordered time-series pairs of open and close prices on the same data the price relatives are then computed as
$$x_{m,t} = \frac{c_{m,t}}{o_{m,t}}. \tag{32}$$
Here one is trying to exploit price relative changes within a trade increment, for example, across a single day from the market opening to the market close ignoring the over-night price changes.

3. *close-to-open*: Here the prices $p_{m,t}$ for the m-th asset are the price changes from the close of the trade period at $t-1$ to the next trade period at time t
$$x_{m,t} = \frac{o_{m,t}}{c_{m,t-1}}. \tag{33}$$
Here one is looking to exploit the change in prices between trade periods where the information cannot yet be fully reflected in trading until the trading commences in the next trade period.

4. *open-to-open*: Here the prices $p_{m,t}$ for the m-th asset are the time-series of opening prices. The price relatives $x_{m,t}$ are then computed from the opening prices

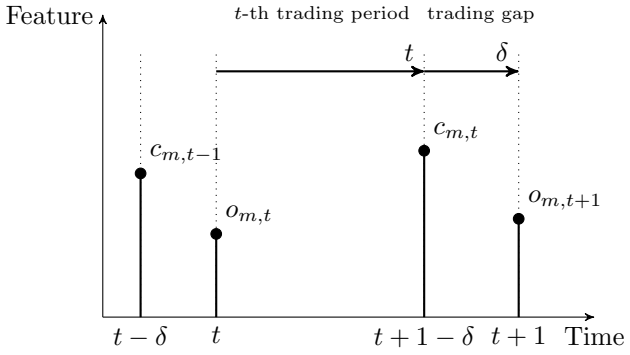

Figure 2: Feature time-series investment period for the t-th time increment showing that the end of the t-th increment does not always have to coincide with the start of the next, here the t+1-th, investment period. The opening price is denote as $o_{m,t}$ and the close price for the period as $c_{m,t}$ for the m-th asset.

$o_{m,t}$

$$x_{m,t} = \frac{o_{m,t}}{o_{m,t-1}}. \tag{34}$$

This is looking for inefficiencies in the prices changes from market opening to market opening.

The important missing component of information is that related to volume (and additional features such as spread, order-imbalance and order-book resilience for intraday data). For example, the opening price is a less reliable price when it has been determined off significantly lower volumes of trading, as compared with a typical closing price. In the case where the closing auction of a given market has more volume than the typical opening auction the relative uncertainties in the prices can be substantial. The typical time increment for a given feature is given in figure (2). We promote the use of a state-detection algorithm and side-information partitioning in order to address these types of concerns. In the context of this work such issues do not change our conclusions.

It is expected that the learning algorithm will still attempt to maximise the long-term wealth given a specific agent generating algorithm for a given feature set. For both daily data and intraday data the feature set that is of most interest to us in this study will be those associated with the "close-to-close" and "close-to-open" price relative features.

4.2. Synthetic Data

The algorithm was tested on four synthetic data cases (SDC) for both active and absolute portfolios. The synthetic data was generated for 10 stocks over 1000 time periods. The price relatives $x_{m,t}$ for each stock at each time period was randomly generated from a lognormal distribution (`lognrnd` function in MATLAB generated using the Mersenne Twister psuedorandom number generator [31] and initialised using a specific seed value), each synthetic data case defines a mean, μ, and variance, v, used to generate the dataset. The mean, $\bar{\mu}$, and standard deviation, $\bar{\sigma}$, of the associated normal distribution is given by :

$$\bar{\mu} = \log\left(\frac{\mu^2}{\sqrt{v+\mu^2}}\right) \tag{35}$$

$$\bar{\sigma} = \sqrt{\log\left(\frac{v}{\mu^2}+1\right)} \tag{36}$$

Table 1 summarises the four synthetic data cases, each case was generated 30 times and initialised with seed values $1, 2, \ldots, 30$ respectively.

1. **Synthetic Data Case 1 (SDC 1):** was generated from a lognormal distribution with a mean price relative, $\mu = 1$, and a variance, $v = 0.0002$, to simulate a stock market where there is no significant increase or decrease in the value of a stock over time.

   The expected outcome is that neither the active portfolio nor the absolute portfolio will be able to learn which stocks it should hold a long position or short position.

2. **Synthetic Data Case 2 (SDC 2):** was generated from a lognormal distribution with a mean price relative, $\mu = 1.001$, and a variance, $v = 0.0002$, to simulate a stock market where the value of the stocks are increasing over time.

   The expected outcome is that the absolute portfolio will learn which stocks to hold a long position on, however the active portfolio will not be able to learn which stocks to hold a short position on as no stocks decrease in value over time.

3. **Synthetic Data Case 3 (SDC 3):** was generated from a lognormal distribution with a random mean price relative, $\mu \geq 1$, assigned to each stock and a variance, $v = 0.0002$. The random means is calculated as follows:

   $$\mu_m = 1 + \max[0, \min(0.0005 + 0.0005\delta, 0.001)] \tag{37}$$

   where δ is a random number generated from a standard normal distribution (using the `randn` function in MATLAB with the Mersenne Twister psuedorandom number generator [31] and initialised using a specific seed value). This simulates a stock market where some stocks are increasing in value and some stocks are decreasing in value over time.

   The expected outcome is that both the active portfolio and the absolute portfolio will learn to hold a long position on the stocks increasing in value over time and hold a short position on the stocks decreasing in value over time, however it is expected that the absolute portfolio will beat the active portfolio due to the growth rate of the stocks increasing in value over time.

Summary of Random Datasets

Dataset	μ	v
SDC 1	1.000	0.0002
SDC 2	1.001	0.0002
SDC 3	random≥ 1	0.0002
SDC 4	mixed	0.0002

Table 1: The means and variances that were chosen when generating the synthetic data sets. The random means for SDC 3 was calculated using Eqn. (37) and the means for SDC 4 was generated as described in section 4.

4. **Synthetic Data Case 4 (SDC 4):** was generated from a lognormal distribution with mixed means assigned to the price relatives, $\mu = 0.999$ was assigned to 3 stocks and $\mu = 1.001$ was assigned to the remaining stocks, and a variance, $v = 0.0002$. This dataset simulates a stock market where the value of some stocks are increasing and the value of some stocks are decreasing.

   The expected outcome is that both the active portfolio and the absolute portfolio will learn to hold a long position on the stocks increasing in value over time and hold a short position on the stocks decreasing in value over time.

4.3. Real Data

The algorithm is tested on four sets of real data, summarised in Table 2, two data sets from the New York Stock Exchange (NYSE) obtained at [32] and two data sets from the Johannesburg Stock Exchange (JSE) obtained at [33].

1. **NYSE Data:** This is described in [32] [12] and contains *close-to-close* price relatives for 36 stocks listed on the New York Stock Exchange from 1962-1984. This is the same data set used by Györfi *et al* in [20, 22] and Cover in [17].

2. **NYSE Merged Data:** This is described in [32][13] and the dataset contains *close-to-close* price relatives data for 23 stocks listed on the New York Stock Exchange from 1962-2006. The data of the 23 stocks during 1962-1984 is identical to the data described above in point 1.

3. **JSE OHLC Data:** This was obtained from Thomson Reuters Tick History (TRTH) [33] and contains daily data for 42 stocks listed on the Johannesburg Stock Exchange (JSE) from 1995-2015 (using RIC chain `0#.JTOPI`), however not all of the 42 stocks were listed in 1995 and the data for these stocks begins at a later time [14]. The data lists the open, high, low and close prices for all of the 42 stocks. This raw data was processed into four datasets containing *close-to-close*, *open-to-close*, *close-to-open* and *open-to-open* price relatives respectively. Splits, mergers[15] and missing data were handled by assigning a price relative of 1 for that day.

Summary of Real Datasets

Data Set	Time Period	# Stocks
NYSE [32]	1962-1984	36
NYSE Merged [32]	1962-2005	23
JSE daily OHLC [33]	1995-2015	42
JSE Intraday [33]	2013	40

Table 2: Description of the real data sets that the algorithm was tested on.

4. **JSE Intraday Data:** The transaction data was obtained from Thomson Reuters Tick History (TRTH) [33] and consisted of top-of-book and transaction updates for 40 stocks listed on the Johannesburg Stock Exchange (JSE) during 2013 in RIC chain `0#.JTOPI`. The transaction data was converted into 5-minute bar data using the trade price and volume weighted averaging. The 5-minute bar-data starts at 9h30 and ends at 16h30 for normal trading days and starts at 9h30 and ends at 11h30 for early close days. A normal trading data on the JSE starts with an opening auction between 8h30 and 9h00, continuous trading takes place between 9h00 and 16h50, and the day ends with a closing auction between 16h50 and 17h00.

5. Implementation

The wealth achieved by the portfolio and the wealth achieved by the agents is determined using Algorithm 1 (OLA). The agent controls $\boldsymbol{H}_{n,t}$, introduced in section 2, used in Algorithm 1 is determined by using Algorithm 2 (PMA), which calls up Algorithm 3 (MTA)[16] to determine the agent controls for each agent. Algorithm 3 updates an agents wealth as described in Eqn. (25). In the experiments 50 agents were used with $K = (1, 2, \ldots, 5)$ and $L = (1, 2, \ldots, 10)$, similar to choice of agents ('experts') used by Györfi *et al* in [20, 22].

All Results and data processing was done in MATLAB. The algorithm was implemented for both the absolute and active case using a MATLAB class that we named *pattern*,

[12] This data set comes from the website of Yoram Singer[32]

[13] This data was original sourced from Yahoo! Finance and was cleaned and prepared by Gábor Gelencsér and made available on his website [32].

[14] The data has an implicit survivorship bias however this does not impact the results of this paper

[15] Splits and mergers were identified as having a $x_{m,t} < 0.7$ and $x_{m,t} > 1.3$ respectively.

[16] Calendar effects are not fully accounted for.

Wealth (S) from Investing in Synthetic Data

Data	Port.	Wealth		Best Agent	
		Best	Avg.	Best	Avg.
SDC 1	Abs.	1.231	0.992	1.806	1.250
	Act.	1.451	1.052	1.753	1.358
SDC 2	Abs.	3.241	2.612	4.654	3.270
	Act.	1.451	1.052	1.753	1.358
SDC 3	Abs.	2.320	1.685	3.091	2.090
	Act.	1.490	1.171	1.782	1.451
SDC 4	Abs.	2.455	1.896	2.931	2.250
	Act.	2.927	2.055	3.270	2.297

Table 3: Wealth achieved by the active and absolute portfolios for 30 runs of each synthetic data case.

a MATLAB class was used instead of function because this allows the algorithm to easily be extended to a more online approach. The *pattern* class was extended to include our recovered version of the Györfi *et al* Nearest Neighbour [22] algorithm so that the running time comparisons in section 6 will be accurate. The Cover *et al* [17] Universal Portfolios algorithm was recovered by creating a MATLAB function that implement the algorithm.

6. Results and Analysis

6.1. Synthetic Data

The algorithm was tested on four synthetic data cases (SDC) to illustrate how the algorithm performs in different types of markets. Table 3 displays the best and average wealth achieved by the active and absolute portfolios for 30 runs of each synthetic data case initialised with seed values $1, 2, \ldots, 30$ respectively.

On all of the datasets the algorithm, when using absolute portfolio, eventually learns the stocks that are increasing in value over time as observed for SDC 2, 3 and 4. Similarly the algorithm, when using active portfolio, eventually learns to hold a long position on the stocks that are increasing in value over time and hold a short position on the stocks that are decreasing in value over time; as observed for SDC 3 and 4. Figures 3, 4, 5 and 6 shows the wealth achieved by the active and absolute portfolios, as well as the wealth achieved by each synthetic stock when randomly generated using an initial seed value of 7.

Tables 4 and 5 displays average p values from the two-sample Kolmogorov-Smirnov tests when comparing the following combinations of the total wealth gained from the portfolio (S_1), the wealth gained from the best agent of the portfolio (S_2) and the wealth gained from the best stock (S_3):

1. $S_2 > S_1$: The alternative hypothesis that the cumulative distribution function (CDF) of the wealth gained from the best agent of the portfolio, S_2, is larger than the CDF of the total wealth gained from the portfolio, S_1, at the 5% significance level.

Average p values of Wealth (S) for Active Portfolios

	Best Agent vs. Tot. Wealth		Best Agent vs. Best Stock		Best Stock vs. Tot. Wealth	
Hyp.	$S_2 > S_1$		$S_2 > S_3$		$S_3 > S_1$	
	$\bar{p}$	$p > \bar{p}$	$\bar{p}$	$p > \bar{p}$	$\bar{p}$	$p > \bar{p}$
SDC 1	0.809	0.172	0.031	0.000	0.654	0.013
SDC 2	0.809	0.172	0.000	0.000	0.873	0.407
SDC 3	0.830	0.172	0.000	0.000	0.904	0.563
SDC 4	0.725	0.013	0.000	0.000	0.622	0.006

Table 4: Comparisons of the average p values of the wealth gained from the active portfolio. The first p value in each column is average p value, of the 30 data sets for each case, using two-sample Kolmogorov-Smirnov tests for the alternative hypotheses (Hyp.). The second p value is obtained from the two-sample Kolmogorov-Smirnov tests for the alternative hypothesis that the cumulative distribution function (CDF) of the p values for the 30 data sets for each case is larger than the CDF of the average p value at the 5% significance level.

2. $S_2 > S_3$: The alternative hypothesis that the CDF of the wealth gained from the best agent of the portfolio, S_2, is larger than the CDF of the wealth gained from the best stock, S_3, at the 5% significance level.

3. $S_3 > S_1$: The alternative hypothesis that the CDF of the wealth gained from the best stock, S_3, is larger than the CDF of the total wealth gained from the portfolio, S_1, at the 5% significance level.

The two-sample Kolmogorov-Smirnov test was chosen because it is a non-parametric test and makes no assumption about the distribution of the datasets.

6.2. NYSE Data

The algorithm was run on the NYSE data set for both absolute and active portfolios on the same pairs of stocks used by Cover in [17] and by Györfi *et al* in [20, 22]. Table 8 shows the wealth achieved by the active and absolute portfolios and is compared to reference results from the literature when using the nearest neighbour strategy (G_{NN}) by Györfi *et al* [20, 22] and the universal portfolio strategy (UP) by Cover [17].

G^*_{NN} denotes our best recovery of the results of the nearest neighbour strategy [20, 22]. The results achieved by from the universal portfolio strategy was identically recovered [17]. The last row of table 8 shows the results of the strategies when running on all 36 NYSE stocks.

The algorithm compares well to the two stocks combinations used by Cover in [17] and by Györfi *et al* in [20, 22]. A surprising result is how the wealth achieved by the portfolio when run over all 36 stocks compares to results by

Average p Values of Wealth Gained (S) from the Absolute Portfolio

	Best Agent vs. Wealth		Best Agent vs. Best Stock		Best Stock vs. Wealth	
Hyp.	$S_2 > S_1$		$S_2 > S_3$		$S_3 > S_1$	
	$\bar{p}$	$p > \bar{p}$	$\bar{p}$	$p > \bar{p}$	$\bar{p}$	$p > \bar{p}$
SDC 1	0.893	0.407	0.000	0.000	0.647	0.013
SDC 2	0.642	0.013	0.000	0.000	0.567	0.001
SDC 3	0.593	0.002	0.000	0.000	0.795	0.100
SDC 4	0.846	0.274	0.000	0.000	0.644	0.006

Table 5: Comparisons of the average p values of the wealth gained from the absolute portfolio. The first p value in each column is the average p value, of the 30 data sets for each case, using two-sample Kolmogorov-Smirnov tests for the alternative hypotheses (Hyp.). The second p value is obtained from the two-sample Kolmogorov-Smirnov tests for the alternative hypothesis that the cumulative distribution function (CDF) of the p values for the 30 data sets for each case is larger than the CDF of the average p value at the 5% significance level.

Comparison of p Values of Wealth Gained (S) from the Active Portfolio

	SDC 1	**SDC 2**	**SDC 3**	**SDC 4**
SDC 1	-	0.9765	0.0669	0
SDC 2	0.9786	-	0.0669	0
SDC 3	0.7928	0.7916	-	0
SDC 4	0.9653	0.9649	0.9535	-

Table 6: Comparison of the average p values from two-sample Kolmogorov-Smirnov tests for the alternative hypothesis that the cumulative distribution function (CDF) of wealth gained from the active portfolio on SDC i is larger than the CDF of wealth gained from the active portfolio on SDC j at the 5% significance level, where i represents the rows and j represents the columns of the table. The p values is the average of 30 comparisons, each comparison using a seed value of $1, 2, \ldots, 30$ respectively.

Comparison of p Values of Wealth Gained (S) from the Absolute Portfolio

	SDC 1	**SDC 2**	**SDC 3**	**SDC 4**
SDC 1	-	0	0	0
SDC 2	1.0000	-	1.0000	0.9997
SDC 3	1.0000	0	-	0.1289
SDC 4	1.0000	0	0.5055	-

Table 7: Comparison of the average p values from two-sample Kolmogorov-Smirnov tests for the alternative hypothesis that the cumulative distribution function (CDF) of wealth gained from the absolute portfolio on SDC i is larger than the CDF of wealth gained from the absolute portfolio on SDC j at the 5% significance level, where i represents the rows and j represents the columns of the table. The p values is the average of 30 comparisons, each comparison using a seed value of $1, 2, \ldots, 30$ respectively.

Synthetic Data Case 1

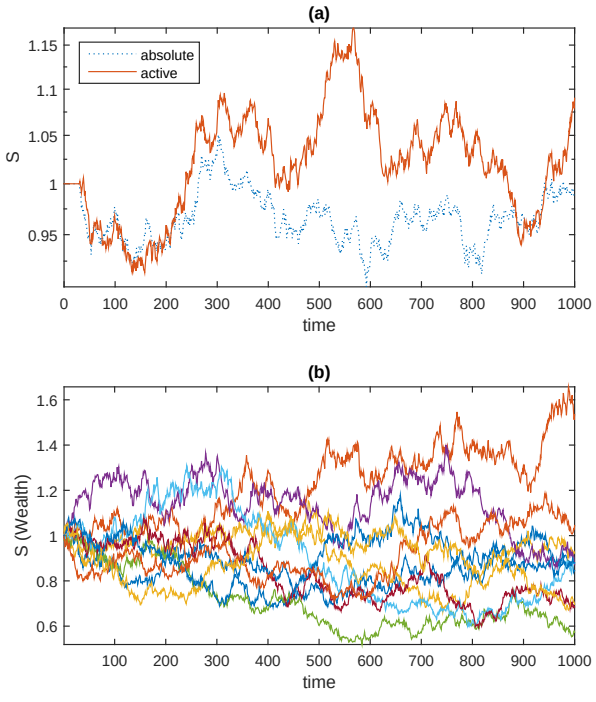

Figure 3: (a) The wealth achieved by the active and absolute portfolios on SDC 1 that consists of a time period of 1000 and 10 stocks. (b) The wealth of each randomly generated stock.

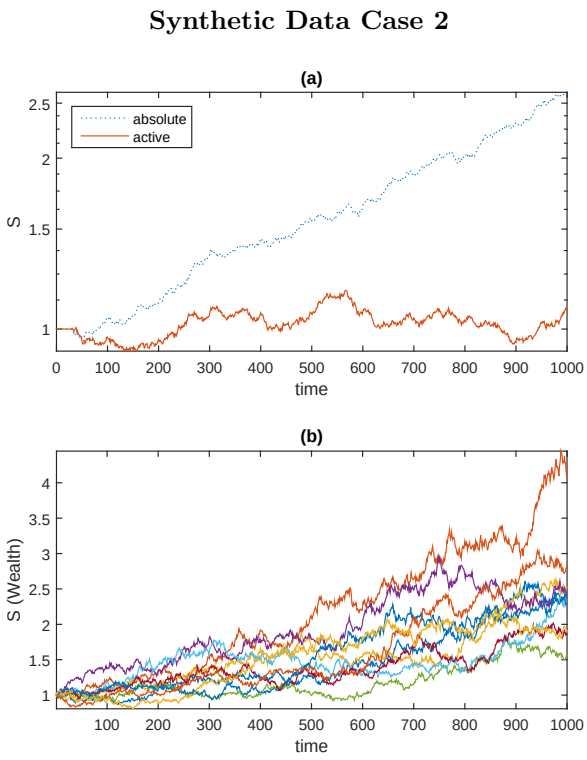

Figure 4: (a) The wealth achieved by the active and absolute portfolios on SDC 2 that consists of a time period of 1000 and 10 stocks. (b) The wealth of each randomly generated stock.

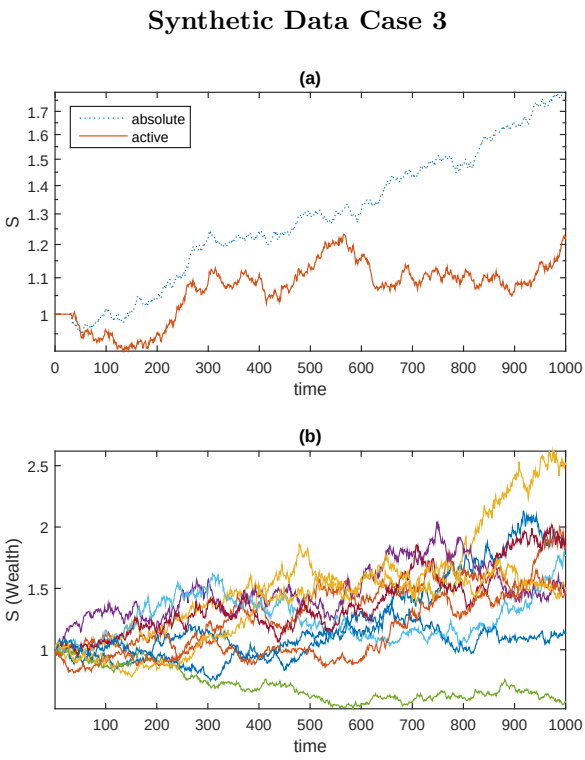

Figure 5: (a) The wealth achieved by the active and absolute portfolios on SDC 3 that consists of a time period of 1000 and 10 stocks.(b) The wealth of each randomly generated stock.

Synthetic Data Case 4

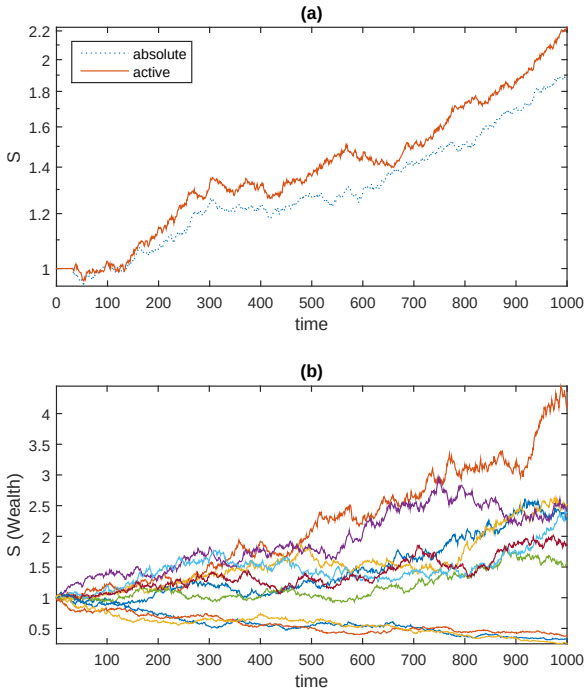

Figure 6: (a) The wealth achieved by the active and absolute portfolios on SDC 4 that consists of a time period of 1000 and 10 stocks. (b) The wealth of each randomly generated stock.

Wealth Achieved by Investing in Various Combinations of Stocks From NYSE Dataset

Stocks	Strat.	Wealth	Best Agent
IROQU KINAR	Abs.	1.02e+12	1.76e+13
	Act.	1.00e+11	4.97e+11
	G_{NN}	1.16e+12	1.44e+13
	G^*_{NN}	1.01e+12	1.63e+13
	UP	38.67	
	Best	8.92	
COMME MEICO	Abs.	3.56e+03	2.61e+04
	Act.	4.28e+01	2.49e+02
	G_{NN}	3.51e+3	3.15e+4
	G^*_{NN}	3.58e+03	2.60e+04
	UP	72.63	
	Best	52.02	
COMME KINAR	Abs.	2.99e+12	3.46e+13
	Act.	1.05e+11	3.75e+11
	G_{NN}	4.78e+12	8.26e+13
	G^*_{NN}	3.09e+12	3.75e+13
	UP	78.47	
	Best	52.02	
IBM COKE	Abs.	7.84e+01	2.74e+02
	Act.	9.70e+00	1.98e+01
	G_{NN}	74.37	296.3
	G^*_{NN}	7.83e+01	2.67e+02
	UP	14.18	
	Best	13.36	
36 STOCKS	Abs.	5.42e+01	1.36e+02
	Act.	5.29e+01	7.13e+01
	G_{NN}	3.3e+11	7.7e+12
	G^*_{NN}	3.43e+11	7.45e+12
	Best	54.14	

Table 8: Comparison of the total wealth achieved from the active (Act.) and absolute (Abs.) portfolios to the wealth achieved from the Györfi *et al* nearest neighbour (G_{NN}), attempted recovery of the Györfi *et al* nearest neighbour (G^*_{NN}), the universal portfolio (UP) and a buy-and-hold of the best stock strategies.

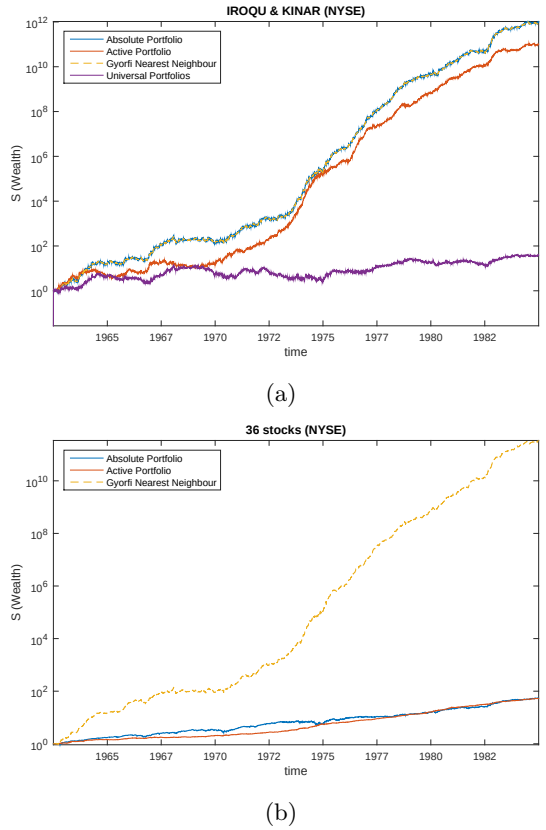

Figure 7: Comparison of the wealth gained from different methods when investing in (a) iroqu and kinar (b) 36 stocks from the NYSE dataset.

Györfi *et al* in [20, 22], this may be due to a loss of accuracy in the quadratic approximation step of the algorithm as the number of stocks increase.

6.3. NYSE Merged Data

The algorithm was run for both absolute and active portfolios on the NYSE Merged dataset on two stock combinations and on all of the 23 stocks in the dataset. The two stock combinations chosen were stocks Commercial Metals and Kin Ark Corp. and stocks IBM and Coca-Cola.

6.4. Daily sampled JSE data

The algorithm was run for both absolute and active portfolios on various sets of two stock combinations, namely stocks AngloGold Ashanti Ltd and Anglo American PLC, stocks Standard Bank Group Ltd and FirstRand Ltd, and stocks Tiger Brands Ltd and Woolworths Holdings Ltd. The algorithm was also run for both absolute and active portfolios on a combination of 10 stocks, 20 stocks and 30 stocks. In each case the date for which the data of a stock starts may be different, the time period for the algorithm therefore starts with the stock that has a

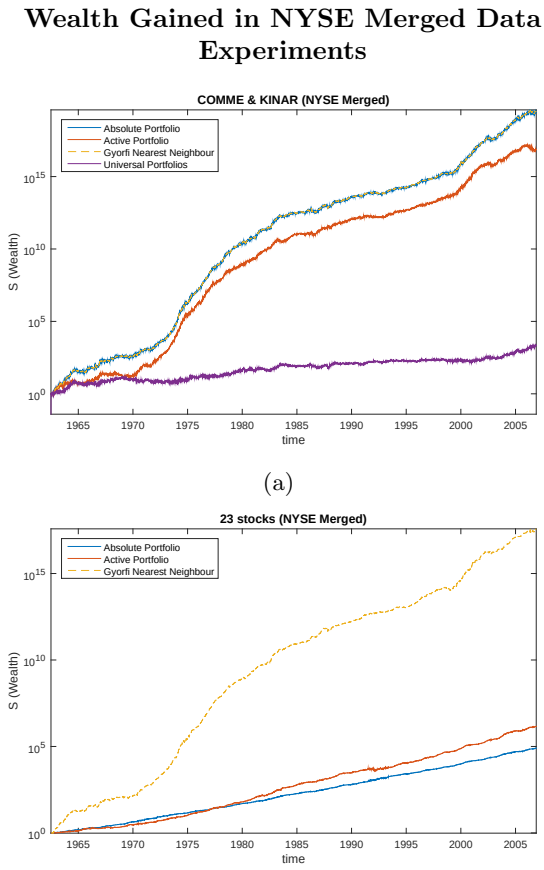

Figure 8: Comparison of the wealth gained from different methods when investing in (a) comme and kinar (b) 23 stocks from the NYSE Merged dataset.

Wealth Achieved by Investing in Various Combinations of Stocks From NYSE Merged Dataset

Stocks	Strat.	Wealth	Best Agent
COMME KINAR	Abs.	3.07e+19	4.37e+20
	Act.	7.36e+16	7.93e+16
	G^*_{NN}	3.19e+19	4.73e+20
	UP	2192.43	
	Best	1344.3	
IBM COKE	Abs.	1.79e+03	5.12e+03
	Act.	7.79e+00	2.63e+01
	G^*_{NN}	1.79e+03	4.75e+03
	UP	229.13	
	Best	365.92	
23 STOCKS	Abs.	8.05e+04	3.34e+05
	Act.	1.45e+06	5.42e+05
	G^*_{NN}	3.68e+17	5.60e+18
	Best	3496.7	

Table 9: Comparison of the total wealth achieved from the active (Act.) and absolute (Abs.) portfolios to the wealth achieved from the attempted recovery of the Györfi *et al* nearest neighbour strategy (G^*_{NN}), the attempted recovery of the universal portfolio strategy (UP^*) and a buy-and-hold strategy of the best stock (Best).

later starting time. The JSE OHLC dataset was processed into four datasets containing *close-to-close*, *open-to-close*, *close-to-open* and *open-to-open* price relatives, the algorithm was run on each of these datasets.

6.5. Intraday JSE data

The algorithm was run for both absolute and active portfolios on various sets of two stock combinations, namely stocks AngloGold Ashanti Ltd and Anglo American PLC, stocks Standard Bank Group Ltd and FirstRand Ltd, stocks Tiger Brands Ltd and Woolworths Holdings Ltd, and stocks MTN Group Ltd and Vodacom Group Ltd. The algorithm was also run for both absolute and active portfolios on the same combination of 10 stocks used on the JSE OHLC Dataset.

6.6. The Impact of Market Frictions

An important criticism of any strategy simulation relates to the need to account for the impact of market frictions, this includes: transaction costs, the cost of the capital for trading, the cost of market access, the cost of regulatory capital for taking risky trading positions, and market impact. These are all required to be included in any estimate of performance slippage for any realistic assessment of the viability of trading activity.

6.6.1. Daily strategy trading frictions

The argument that the zero-cost low frequency (daily traded) strategies are viable, even when unleveraged, is

Wealth Achieved by Investing in Various Combinations of Stocks From JSE OHLC Dataset (close-to-close)

Stocks	Strat.	Wealth	Best Agent
ANGJ AGLJ	Abs.	4.05	13.10
	Act.	1.28	3.40
	G^*_{NN}	4.02	13.27
	UP	2.52	
	Best	3.61	
SBKJ FSRJ	Abs.	55.53	187.62
	Act.	7.77	13.40
	G^*_{NN}	55.60	189.08
	UP	18.17	
	Best	21.09	
TBSJ WHLJ	Abs.	7.07	19.97
	Act.	0.49	1.80
	G^*_{NN}	7.07	19.49
	UP	8.24	
	Best	8.97	
10 STOCKS	Abs.	68.49	135.03
	Act.	9.28	12.84
	G^*_{NN}	194.76	854.62
	Best	89.72	
20 STOCKS	Abs.	16.20	63.18
	Act.	9.52	12.11
	G^*_{NN}	98.84	330.37
	Best	89.72	
30 STOCKS	Abs.	20.48	58.27
	Act.	7.34	6.24
	G^*_{NN}	124.91	590.51
	Best	85.03	

Table 10: The total wealth achieved by the active (Act.) and absolute (Abs.) portfolios compared to the wealth achieved from the attempted recovery of the Györfi *et al* nearest neighbour strategy (G^*_{NN}), the attempted recovery of the universal portfolio strategy (UP^*) and a buy-and-hold strategy of the best stock (Best) on the *close-to-close* dataset.

Wealth Achieved by Investing in Various Combinations of Stocks From JSE OHLC Dataset (close-to-open)

Stocks	Strat.	Wealth	Best Agent
ANGJ AGLJ	Abs.	1.24	2.60
	Act.	1.04	2.82
	G^*_{NN}	1.24	2.56
	UP	0.97	
	Best	2.66	
SBKJ FSRJ	Abs.	12.35	25.13
	Act.	1.57	3.06
	G^*_{NN}	12.31	25.14
	UP	7.89	
	Best	12.47	
TBSJ WHLJ	Abs.	4.38	9.00
	Act.	0.79	1.88
	G^*_{NN}	4.46	8.83
	UP	4.89	
	Best	4.98	
10 STOCKS	Abs.	7.11	11.43
	Act.	2.68	3.39
	G^*_{NN}	13.45	30.53
	Best	56.36	
20 STOCKS	Abs.	8.82	16.87
	Act.	5.28	5.24
	G^*_{NN}	11.00	27.01
	Best	56.36	
30 STOCKS	Abs.	8.07	15.80
	Act.	5.66	5.50
	G^*_{NN}	22.53	54.80
	Best	51.02	

Table 11: The total wealth achieved by the active (Act.) and absolute (Abs.) portfolios compared to the wealth achieved from the attempted recovery of the Györfi *et al* nearest neighbour strategy (G^*_{NN}), the attempted recovery of the universal portfolio strategy (UP^*) and a buy-and-hold strategy of the best stock (Best) on the *close-to-open* dataset.

Wealth Achieved by Investing in Various Combinations of Stocks From JSE OHLC Dataset (open-to-close)

Stocks	Strat.	Wealth	Best Agent
ANGJ AGLJ	Abs.	97.75	395.58
	Act.	50.21	91.82
	G^*_{NN}	97.53	398.74
	UP	3.41	
	Best	4.68	
SBKJ FSRJ	Abs.	2.53	5.12
	Act.	1.04	1.80
	G^*_{NN}	2.53	5.00
	UP	2.66	
	Best	2.97	
TBSJ WHLJ	Abs.	5.59	12.73
	Act.	5.00	7.19
	G^*_{NN}	5.59	12.75
	UP	1.94	
	Best	1.89	
10 STOCKS	Abs.	15.90	33.40
	Act.	15.21	19.86
	G^*_{NN}	101.80	415.97
	Best	169.83	
20 STOCKS	Abs.	4.80	8.88
	Act.	6.80	6.07
	G^*_{NN}	30.26	104.46
	Best	169.83	
30 STOCKS	Abs.	5.16	11.98
	Act.	4.58	5.57
	G^*_{NN}	50.17	317.92
	Best	1278.83	

Table 12: The total wealth achieved by the active (Act.) and absolute (Abs.) portfolios compared to the wealth achieved from the attempted recovery of the Györfi *et al* nearest neighbour strategy (G^*_{NN}), the attempted recovery of the universal portfolio strategy (UP^*) and a buy-and-hold strategy of the best stock (Best) on the *open-to-close* dataset.

Wealth Achieved by Investing in Various Combinations of Stocks From JSE OHLC Dataset (open-to-open)

Stocks	Strat.	Wealth	Best Agent
ANGJ AGLJ	Abs.	5.11	26.65
	Act.	1.43	9.10
	G^*_{NN}	5.07	25.66
	UP	2.28	
	Best	3.18	
SBKJ FSRJ	Abs.	113.43	412.75
	Act.	13.97	18.87
	G^*_{NN}	113.02	415.45
	UP	20.29	
	Best	20.55	
TBSJ WHLJ	Abs.	45.30	266.17
	Act.	5.52	28.54
	G^*_{NN}	45.80	273.44
	UP	8.81	
	Best	9.24	
10 STOCKS	Abs.	60.62	342.40
	Act.	12.03	20.57
	G^*_{NN}	254.00	2292.47
	Best	69.49	
20 STOCKS	Abs.	14.73	73.61
	Act.	9.13	14.76
	G^*_{NN}	44.88	192.99
	Best	69.49	
30 STOCKS	Abs.	21.54	73.88
	Act.	11.48	18.89
	G^*_{NN}	161.16	1069.81
	Best	131.89	

Table 13: The total wealth achieved by the active (Act.) and absolute (Abs.) portfolios compared to the wealth achieved from the attempted recovery of the Györfi *et al* nearest neighbour strategy (G^*_{NN}), the attempted recovery of the universal portfolio strategy (UP^*) and a buy-and-hold strategy of the best stock (Best) on the *open-to-open* dataset.

Wealth Gained close-to-close JSE OHLC Data

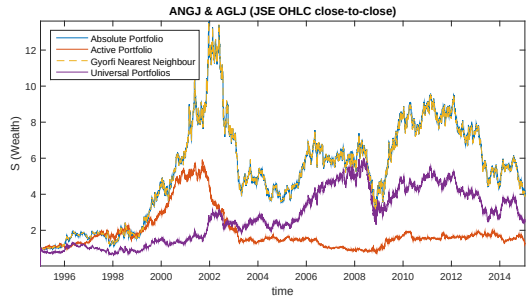

(a)

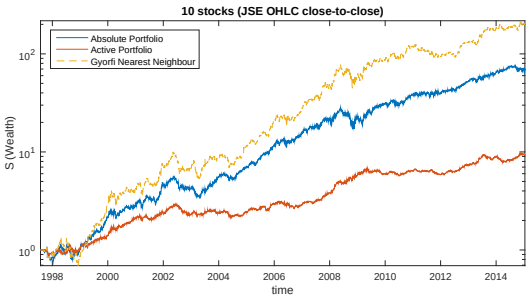

(b)

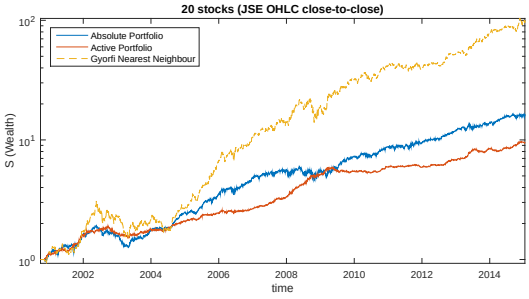

(c)

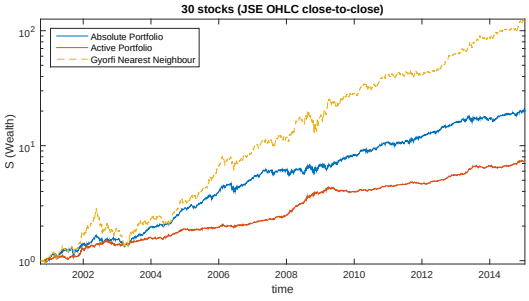

(d)

Figure 9: Comparison of the wealth gained from different methods when investing in (a) ANGJ and AGLJ (b) 10 stocks (c) 20 stocks (d) 30 stocks from the JSE OHLC *close-close* dataset. This does not account for price-impacts and frictions, nor for the need to approximate an expected close price just prior to market close as one solves for the portfolio controls, there will always be a difference between the controls solved for just prior to market close and those required once the market has closed and the official closing prices printed.

Wealth from Different JSE OHLC Data Sets

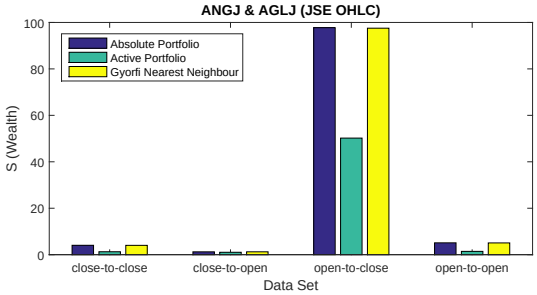

(a)

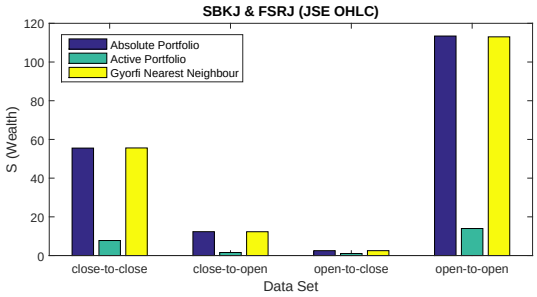

(b)

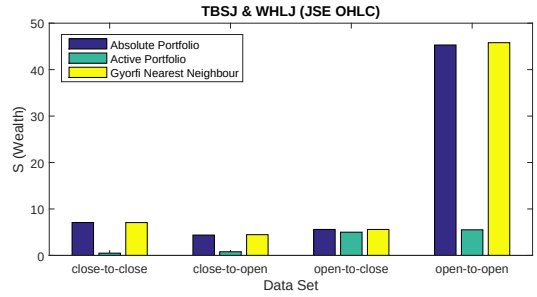

(c)

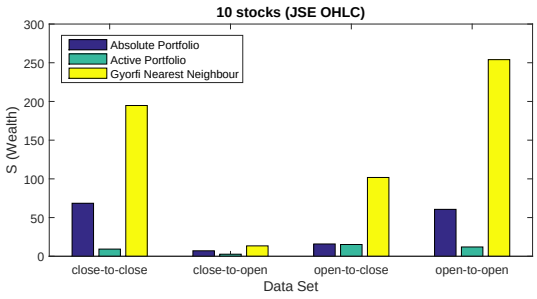

(d)

Figure 10: Comparison of wealth achieved from the absolute portfolio, active portfolio and Györfi nearest neighbour porfolio on the(a)*close-to-close* (b) *open-to-close* (c) *close-to-open* and (d) *open-to-open* JSE OHLC datasets. Here we find that there is no particular combination of OHLC data for which there is a systematic preference, e.g. *close-to-close*, the case of considering the close price change from one day end to another is not systematically more profitable than other combinations of data times. These tests do consider the reality of trading prior to a time point, for example market close, one cannot *a-priori* know what the close price will be, this has to be approximated. This excludes price-impact effects.

Wealth from Stock-Pairs of JSE Intraday Data

Stocks	Strat.	Wealth	Best Agent
ANGJ AGLJ	Abs.	1.38	3.33
	Act.	2.16	5.01
	G^*_{NN}	1.36	3.21
	UP	0.66	
	Best	0.87	
SBKJ FSRJ	Abs.	1.82	2.19
	Act.	2.02	2.01
	G^*_{NN}	1.82	2.15
	UP	1.08	
	Best	1.08	
TBSJ WHLJ	Abs.	1.95	3.06
	Act.	2.24	3.29
	G^*_{NN}	1.95	2.99
	UP	0.91	
	Best	0.99	
MTNJ VODJ	Abs.	2.13	3.55
	Act.	2.13	3.25
	G^*_{NN}	2.13	3.57
	UP	1.12	
	Best	1.18	
10 STOCKS	Abs.	1.89	2.79
	Act.	3.86	5.40
	G^*_{NN}	3.95	14.05
	Best	1.93	
20 STOCKS	Abs.	1.74	2.35
	Act.	3.42	4.29
	G^*_{NN}	5.68	21.62
	Best	1.93	
30 STOCKS	Abs.	1.67	2.16
	Act.	3.07	3.76
	G^*_{NN}	6.03	12.61
	Best	1.93	

Table 14: The total wealth achieved by the active (Act.) and absolute (Abs.) portfolios compared to the wealth achieved from the attempted recovery of the Györfi *et al* nearest neighbour strategy (G^*_{NN}), the attempted recovery of the universal portfolio strategy (UP^*) and a buy-and-hold strategy of the best stock (Best) on the JSE Intraday Dataset.

Wealth from JSE Intraday Data Experiments

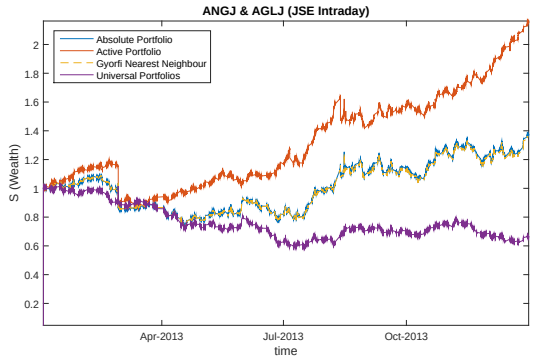

(a)

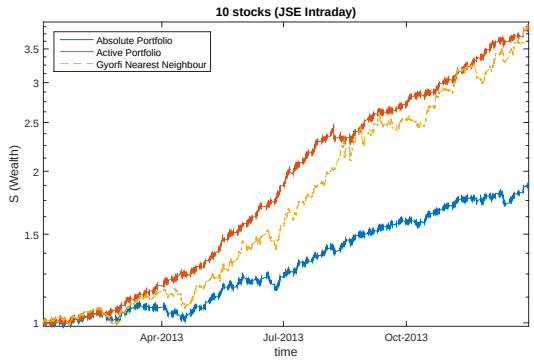

(b)

Figure 11: Comparison of the wealth gained from different methods when investing in (a) ANGJ and AGLJ (b) 10 stocks from the JSE Intraday dataset.

JSE Intraday Data with cluster defined agents

		RESI	INDI	FINI
Act.	Best Agent	9.1018	3.4470	3.6698
	Total Wealth	4.6368		
Abs.	Best Agent	5.6655	2.6341	2.6059
	Total Wealth	2.2093		

Table 15: Wealth achieved by active and absolute portfolios when using economic sectors as clusters. Using three clusters increases the number of competing agents by a factor of 3, from 50 to 150. The inclusion of a larger set of agents increase the out-of-sample wealth performance of the two strategies. See Tables Appendix A.2.2, Appendix A.2.1, and Appendix A.2.3 for the particular stocks in each sectors, Resources (RESI), Financials (FINI) and Industrials (INDI), respectively. The total wealth includes the relative competition between agents defined by the three economic sectors. The sector specific wealth is given the best performing agents from each cluster groups of stocks.

20 JSE Stocks

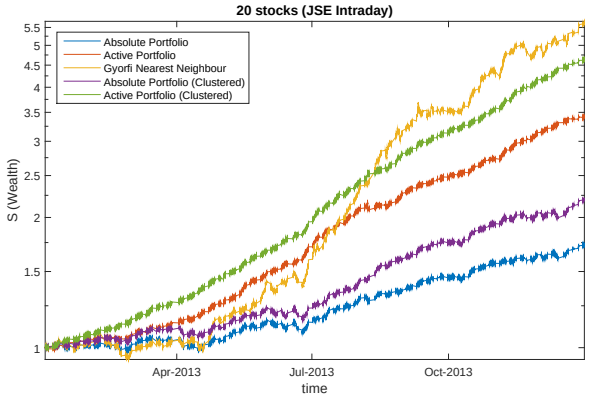

Figure 12: Comparison of the wealth gained from different methods when investing in 20 stocks from the JSE Intraday dataset, the plot includes the results of using clusters on the stocks. It is important to note that the clustered portfolios have 150 agents and the portfolios without clusters have 50 agents.

JSE Intraday Data Running Times

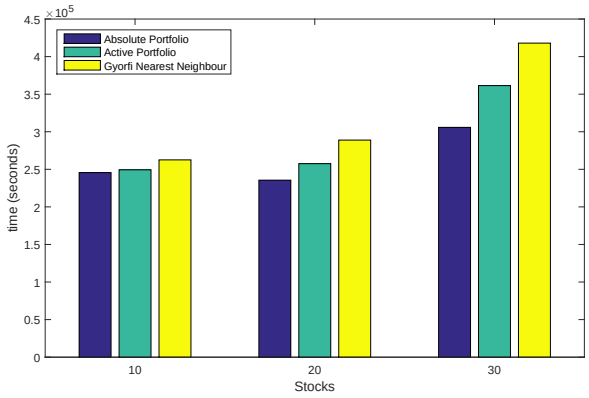

Figure 13: Running time of the portfolios in seconds of the different strategies. This demonstrates the speed advantage of using the analytic quadratic approximation as compared to numerically solving the log-optimal constrained optimization at each time-step for each agent combination. As expected the fully invested analytic solution is fastest, the zero-cost portfolio next, because of the additional leverage constraint, and the slowest the algorithm that required the numerical solution of the optimization.

based on the 7 years of history [17]. Consider Table 10 for the active case for the Top 10 JSE stocks (see appendix Appendix A). Here we would argue for 15bps of daily profit before costs (from Table 10 using the accumulated daily wealth of 9.53). We consider the strategy that trades close of the one day to the close of the next day (close-to-close).

This is considered in order to take into account liquidity effects. The closing auction is the most liquid time to trade on the JSE. It is unlikely that one would be able to achieve low slippage trading near the daily market opening. We consider the combination of cost of capital (the borrowing costs required to source trading capital and cost of regulatory capital) and a small penalty for slippage due to the differences between the realized closing prices and the estimated closing prices that the algorithm would require in order to estimate the portfolio controls [18] as 10bps per day. In practise it should also be noted that such a trading strategy can be converted to one that trades in equity swaps, so-called contracts-for-difference (CFDs), this would convert the uncertainty about slippage into an up-front fee and allow for excellent implementation of the required model positions with a known cost and no meaningful liquidity concerns. If the daily strategy was implemented with these types of delta-one instruments our estimates of slippage can then be considered conservative.

We argue that we can realistically earn a modest 5bps of unleveraged self-funded trading profit per day, or an annual return of 12% of unleveraged profit-and-loss.

6.6.2. Intraday strategy trading frictions

We assume: 1.) a daily slippage of 50bps for the (self-financing) zero-cost statistical arbitrage strategy, 2.) borrowing costs on the capital required for trading over the year to be 10%, and 3.) that the strategy we denote as the active strategy generated a 4.63 wealth gained over a year of trading (see Table 15). Putting these together we argue for an upper limit on the profit, even when unleveraged, to be a return of 20% for 250 days of trading[19].

The strategy turn-over is important in realistic assessments of profitability for intraday trading. We try to account for this in our indicative costing of the slippage by assuming that we have 100% turn-over of inventory at each trading period with a consistent cost of 0.55bps (0.0055%), per trading period[20] with an additional 4bps to give an indicative slippage of 50bps for the intraday trading per day.

7. Conclusion

In prior work it has been shown that in South African financial markets persistence and long-memory are generic [2]. This paper adds to our knowledge of the South African market by showing that in addition to evidence supporting long-memory processes, price processes have patterns that are exploitable in a straight-forward manner.

We provide a simple portfolio value based learning algorithm, a multi-manager in the language of asset management, that selects an over-all portfolio with weights $\boldsymbol{b}$ by considering a selection of N different strategies $\boldsymbol{H}_n$ with their underlying portfolio weights being constructed for underlying strategies that are enumerated over a variety of combinations of time-series patterns, time-scales, clusters and partitions. This is considered in the context of universally consistent strategies [17, 22] but with an extension to directly consider self-financing zero-cost quantitative trading strategies - what we call the active portfolio.

When applying the algorithms to real daily test data, it compares well to results from our implementation of algorithms from the literature [17, 22] and actual results from the literature from the New York Stock Exchange (NYSE) dataset (see Tables 8 and 9).

The active version of the algorithm, when applied to intraday data from the Johannesburg Stock Exchange (JSE), is shown to have performed well in comparison to the best stock, and compares favourably with methods from prior work [17, 22] (see 14). We show that on the Johannesburg Stock Exchange data the algorithms can learn trends and patterns and enhance out-of-sample wealth accumulation for both daily and intraday applications (see Tables 10 and 14). This is demonstrated on both low frequency data, daily sampled data, and higher frequency data, intraday uniformly sampled transaction data.

We have shown that there is an advantage to include agents that are clustered on stock economic sector classifications (see Table 15); this increases the number of agents (or experts) considered by the learning algorithm through including sector membership into the resource, financial and industrial stocks sectors, which in turn boosts the out-of-sample performance. This suggests that combining more sophisticated clustering algorithms [27] with

[17] Although this may be considered short in the context of many academic studies, this is of the order of the time-scale of the business cycle so we considered this realistic.

[18] This can in practice be carried out during the closing auction, just prior to the market close, by estimating online, the equilibrium price that could be the result of maximizing execution volume for the lowest surplus when the market is cleared at market close, bearing in mind that there is a short randomization period at the end of the auction that needs to be accounted for. It is fairly straight-forward to estimate sufficiently reasonable market clearing prices. If this is not considered realistic, one can then merely consider the trading to have occurred during a post market close period (such as that found on the LSE and JSE) where one can transact at the market close price but at lower volumes.

[19] The strategy generates 62bps per day, the slippage is 50bps, leaving 12bps, less the 5bps for the cost of capital, then leaves 7bps to accumulate as profit-and-loss per trading day.

[20] If we considered trading to be for a 7 hour period starting a half hour after market opening, and stopping 15 minutes before the market closing auction (using the JSE market times), this then leads to $84 = 12 \times 7$ 5-minute trading periods across the day, this would give the worst case scenario of 8400% turn-over per day, at a transaction cost of 0.55bps per trading period, we argue that this then leads to a transaction cost of 46bps per day, additional frictions of 4bps are added to this to get the over-all daily slippage estimate of 50bps per day.

machine learning can be advantageous in the domain of quantitative trading.

The pay-off between computational performance and wealth accumulation can be seen by considering the increased duration of the simulation as one increases from 10 stocks, to 20 stocks through to 30 stocks in Figure 13. The commensurate loss in performance can be seen in Table 14. For example, the 20 stock simulation generated a wealth of 1.74 for the absolute portfolio and 5.68 for the Gyorfi *et al* nearest-neighbour strategy with the absolute portfolio being almost 5×10^4 seconds faster (or 18% faster). For intraday statistical arbitrage problems for quantitative trading with many (50>) assets, computational delays can lead to lags between information arrival and order-execution that can negatively impact a strategies profit-and-loss performance.

We have shown that in the daily dataset for the Johannesburg Stock Exchange, when considering open, high, low and close price data, there is an advantage when considering strategies that relate to the patterns arising across closing price to closing price data (see Table 10). It is difficult to profitable trade the market opening price to the market closing price as intraday dynamics seems to become important and one tends to incur significant market frictions associated with poor market liquidity near the market opening. This provides evidence that one can in principle beat the best stock (or the money market account in the case of the self-financing strategy) as pattern persistence is sufficiently robust in the markets considered.

Towards addressing the key criticism related to correctly estimating the impact of market frictions, in Section 6.6 we give our estimates for the impact of market frictions on both the intraday strategy, where an estimated annual return of upto 20% for the unleveraged self-financing strategy, and the daily strategy, trading the closing price of the market from one day to the next, at an annual return of upto 12% (see Section 6.6). Using this we argue that the self-financing zero-cost portfolio strategy can be considered tractable both intraday and across days, that after a reasonable estimates of costs, one is still able to learn how to exploit patterns the recur in the financial times-series data considered in this study. It should also be noted that for optimised intraday trading the event-time paradigm should be implemented rather than the calendar-time approach that was used for simplicity in the experiments in this paper. This is fairly straight-forward to implement using equal volume buckets and online down-sampling transaction data to a time-series of volume-weighted average prices for equal volume buckets [35].

The aim here is to show that there are repeated patterns that can be exploited on both daily samples and intraday time-scales. We do not claim that being able to exploit such patterns is necessarily profitable as a commercial enterprise, what we are claiming is that structure does exist in the financial market time-series that is indicative of existence of repeated structures that emerge and change through time, but after reasonable costs can be considered a riskless profit, or at least be considered a signature of the ability to generate systematic profits from patterns in financial time-series data.

We do not know whether there is a finite state representation of the system that could be used to generate the observed time-series dynamics. We have evidence for non-linear structure in the time-series data, by providing a simplistic algorithm that can exploit structure in time-series data, when it exists, and we know that the algorithm would behave quite differently for random data. To show that this is indicative of some finite state representation would require online state-detection, either via some type of cluster methodology [28], or via some sort of state-space reconstruction algorithm following the methods of deterministic chaos [5, 6, 7]. This paper makes no statement about the existence of a finite and sufficiently stable finite state representation.

The other criticisms could relate to both barriers to entry to reasonably cost effect market access, as well as the scalability of these types of strategies due to stock liquidity. In terms of the prior, many proprietary trading structures within hedge-funds and banks would have very low transaction costs due to bulk trading activities - hence we consider our daily transaction costs of 50bps as onerous but realistic. In terms of the liquidity concerns, we have limited ourselves, in the Johannesburg Stock Exchange data set, to collections of the 10 and 20 most liquid stocks. These stocks can be traded in meaningful volumes.

We could speculate that it is the buying and selling patterns of large institutional mutual funds, or capital flows of large institutional participants in capital markets, that create key feed-backs which generate persistence in patterns of price dynamics. Realistically there are a variety of potential candidate sources of top-down and bottom-up feedbacks within a system as complex and adaptive as the financial market systems; these could provide various mechanisms that can balance disorder with order in a meta-stable configuration of states over various time-scales [3]. We argue that fairly naive computational learning agents can generate wealth within the system without special insights or understanding of the system itself.

8. Acknowledgements

TG would like to thank AIMS South Africa for their support and hospitality at Muizenberg. The authors would like to thank Diane Wilcox for conceptual contributions, Turgay Celik for discussions and ideas relating to algorithm testing, Dieter Hendricks and Raphael Nkomo for various conversations relating the quantitative trading and machine learning for trading. This work was in part supported by NRF grant number 89250. The conclusions herein are due to the authors and the NRF accepts no liability in this regard.

References

References

[1] B. Arthur (1995) Complexity in Economics and Financial Markets, *Complexity* **1** (1), 20–25.

[2] D. Wilcox & T. Gebbie (2008) Periodicity and scaling of eigenmodes in an emerging market, *International Journal of Theoretical and Applied Finance* **11** (7), 739–760.

[3] D. Wilcox & T. Gebbie (2014) Hierarchical Causality in Financial Economics (August 24, 2014). Available at SSRN: http://ssrn.com/abstract=2544327.

[4] J. P. Crutchfield, (2011), Between order and chaos, Nature Physics, Vol. 8, January 2012, 17-23

[5] N. H. Packard, J. P. Crutchfield, J. D. Farmer and R. S. Shaw (1980), Geometry from a time series, Phys. Rev. Lett. 45 7126

[6] M. Casdagli, S. Eubank, J. D. Farmer and J. Gibson (1991) State space reconstruction in the presence of noise Physica D 51 5298

[7] T. Sauer, J. Yorke and M. Casdagli (1991), Embeddology, J. Stat. Phys. 65 579616

[8] E. Bradley (1998) Nonlinear Time-Series Analysis, CU-CS-867-98, Computer Science Technical Reports. Paper 815. `http://scholar.colorado.edu/csci_techreports/815`

[9] C. R. Shalizi, J. P. Crutchfield (2001) Computational mechanics: Patterns and prediction, structure and simplicity, J. Stat. Phys. 104:817-879 http://arxiv.org/pdf/cond-mat/9907176v2.pdf

[10] Shalizi CR, Shalizi KL, Crutchfield JP (2002) Pattern Discovery in Time Series, Part I: Theory, Algorithm, Analysis, and Convergence. http://arXiv.org/abs/cs.LG/0210025.

[11] A. Lo, H. Mamaysky, J. Wang, (2000) Foundations of Technical Analysis:Computational Algorithms, Statistical Inference, and Empirical Implementation, The Journal of Finance, Vol. LV, No. 4., Aug. 2000.

[12] R. Jarrow, M. Teo, Y-K. Tse, M. Warachka (2012), An improved test for statistical arbitrage, Journal of Financial Markets, v.15, no.1, 2012 Feb, p.47(34)

[13] J. L. Kelly (1956) A new interpretation of information rate, *Bell. Syst. Tech.* **35**, 917-926.

[14] P. H. Algoet & T. M. Cover (1988) Asymptotic optimality and symptotic equipartition properties of log-optimum investments, *Annals of Probability*, **16** (2), 876–808.

[15] A. Blum & A. Kalai (1999) Universal Portfolios With and Without Transaction Costs, *Machine Learning* **35** (3), 193–205.

[16] A. Bordin, R. El-Yaniv & V. Gogan (2004) "Can we learn to beat the best stock", *Journal of Artificial Intelligence Research* **21** , 579–594.

[17] T. M. Cover (1991) Universal Portfolios,*Math. Finance.* **1** (1), 1–29.

[18] T. M. Cover & E. Ordentlich (1996) Universal Portfolios with Side Information, *IEEE Transactions on Information Theory* **42** (2), 348–363.

[19] J. E. Cross & A. R. Barron (2003) Efficient Universal Portfolios for Past-Dependent Target Classes, *Mathematical Finance* **13** (2), 245–276.999

[20] L. Györfi, A. Urban, I. Vajda (2007) Kernel-based semi-log-optimal empirical portfolio selection strategies, *International Journal of Theoretical and Applied Finance* **10** (3), 505–516.

[21] L. Györfi, F. Udina, H. Walk (2007) Experiments on universal portfolio selection using data from real markets. Working paper `http://tukey.upf.es/papers/NNexp.pdf`

[22] L. Györfi, F. Udina, H. Walk (2008) Nonparametric nearest neighbor based empirical portfolio selection strategies, Statistics & Risk Modeling, De Gruyter, 26(2), 145-157

[23] D. P. Helmbold, R. E. Schapire, Y. Singer & M. K. Warmuth (2007) On-line portfolio selection using multiplicative updates, *Mathematical Finance* **8** (4), 325–347.

[24] W. Lee (2000) Theory and Methodology of Tactical Asset Allocation, Frank Frabozzi Associates.

[25] R. Merton (1995) Continous Time Finance, Chapter 4, Blackwell.

[26] A. F. Perold & W. F. Sharpe (1988) Dynamic Strategies for Asset Allocation, *Financial Analysts Journal* **44** (1), 16–27.

[27] D. Hendricks, D. Wilcox & T. Gebbie (2016) High-speed detection of emergent market clustering via an unsupervised parallel genetic algorithm. S Afr J Sci. 2016;112 (1/2), Art. 2014-0340 http://dx.doi.org/10.17159/sajs.2016/20140340

[28] D. Hendricks, T. Gebbie & D. Wilcox (2016) Detecting intraday financial market states using temporal clusterings Quantitative Finance, doi:10.1080/14697688.2016.1171378

[29] Industry Classification Benchmark `http://www.icbenchmark.com/ICBDocs/Structure_Defs_English.pdf`

[30] I. Vajda (2006) Analysis of semi-log-optimal investment strategies, *Prague Stochastics* eds. M. Huskova and M. Janzura (Matfyz press, Prague, 2006), pp. 719727.

[31] M. Matsumoto & T. Nishimura (1998) Mersenne Twister: A 623-dimensionally equidistributed uniform pseudo-random number generator, *ACM Transactions on Modeling and Computer Simulation* **8**, 3–30.

[32] G. Gelencsér & G. Ottucsák (2006) NYSE data sets at the log-optimal portfolio homepage, `http://www.szit.bme.hu/~oti/portfolio`

[33] Thomson Reuters (2015), Thomson Reuters Tick History, `https://tickhistory.thomsonreuters.com`.

[34] N. Jegadeesh (1990) Evidence of predictable behavior of security returns. The Journal of Finance, 45(3), 881-898.

[35] D. Easley, M.M. López de Prado, and M. O'Hara. The volume clock: Insights into the high-frequency paradigm (digest summary). Journal of Portfolio Management, 39(1):19(29), 2012.

Appendix A. Ticker lists

Appendix A.1. NYSE

Name	NYSE	NYSE Merged
3M Company	✓	✓
Alcoa	✓	✓
Altria Group	✓	✓
Arco	✓	
Coca Cola	✓	✓
Commercial Metals	✓	✓
Dow Chemical Company	✓	✓
DuPont	✓	✓
Eastman Kodak	✓	✓
Espey	✓	
Exxon Mobil	✓	
Fischback	✓	
Ford Motor Company	✓	✓
Fortune Brands	✓	✓
General Electric	✓	✓
General Motors	✓	✓
Gran Tierra Energy Inc.	✓	
Gulf Oill	✓	
Hewlett-Packard	✓	✓
IBM	✓	✓
Ingersoll-Rand Plc	✓	✓
Iroquois Ltd.	✓	
Johnson & Johnson	✓	✓
Kimberly-Clark Corp	✓	✓
Lukens	✓	
Mei Corp.	✓	
Merck & Company	✓	✓
Mobil	✓	
Kin Ark Corp.	✓	✓
Pillsbury	✓	
Procter & Gamble	✓	✓
Schlumberger N.V.	✓	✓
Sears	✓	
Sherwin-Williams	✓	✓
Texaco	✓	
Wyeth	✓	✓

Appendix A.2. JSE intraday and daily data

In the following tables $\checkmark^*$ represents stocks that are included in the JSE daily dataset grouping and not the JSE intraday dataset grouping, and $\checkmark^+$ represents stocks that are include in the JSE intraday dataset grouping and not the JSE daily dataset grouping.

Appendix A.2.1. Financials / JSE-FINI (J212)

Name	RIC	Stock grouping		
Financial (JSE-FINI)	J212	10	20	30
Standard Bank Grp. Ltd	SBKJ.J	✓	✓	✓
Firstrand Ltd	FSRJ.J		✓	✓
Absa Group Ltd	ASAJ.J		$\checkmark^*$	$\checkmark^*$
Old Mutual Plc	OMLJ.J		✓	✓
Nedbank Group Ltd	NEDJ.J		✓	✓
Sanlam Ltd	SLMJ.J		✓	✓
Investec Plc	INPJ.J			
RMB Holdings Ltd	RMHJ.J			
Growthpoint Prop Ltd	GRTJ.J			✓
African Bank Inv. Ltd	ABLJ.J			
Capital Shop Cnt. Grp Plc	CSOJ.J			
Reinet Investments Sca	REIJ.J			
Redefine Properties Ltd	RDFJ.J			
Discovery Holdings Ltd	DSYJ.J			✓
Liberty Holdings Ltd	LBHJ.J			

Appendix A.2.2. Resources / JSE-RESI (J201)

Name	RIC	Stock grouping		
Resources (JSE-RESI)	J201	10	20	30
BHP Billiton PLC	BILJ.J	✓	✓	✓
Anglo American PLC	AGLJ.J	✓	✓	✓
Sasol LTD	SOLJ.J	✓	✓	✓
Anglo Platinum Ltd	AMSJ.J		✓	✓
Impala Platinum Hld. Ltd	IMPJ.J	✓	✓	✓
Anglogold Ashanti Ltd	ANGJ.J	✓	✓	✓
Gold Fileds Ltd	GFIJ.J		✓	✓
Exxaro Resources Ltd	EXXJ.J			
African Rainbow Minerals	ARIJ.J			✓
Lonmin PLC	LONJ.J			
Harmony G M Co Ltd	HARJ.J			
Assore Ltd	ASRJ.J			
Northam Platinum Ltd	NHMJ.J			
Optimum Coal Hldgs Ltd	OPTJ.J			
Merafe Resources Ltd	MRFJ.J			
Petmin Ltd	PETJ. J			
Wesizwe Platinum Ltd	WEZJ.J			
Sentula Mining Ltd	SNUJ.J			
DRDGold Ltd	DRDJ.J			
Simmer and Jack Mines	SIMJ.J			

Appendix A.2.3. Industrials/JSE-INDI (J211)

Name	RIC	Stock grouping		
Industrials (JSE-INDI)	J211	10	20	30
SABMiller Plc	SABJ.J	✓	✓	✓
MTN Group Ltd	MTNJ.J	✓	✓	✓
Comp. Fin Richemont	CFRJ.J	✓	✓	✓
Naspers Ltd	NPNnJ.J	✓	✓	✓
Kumba Iron Ore Ltd	KIOJ.J			
Vodacom Group Ltd	VODJ.J			
Bidvest Ltd	BVTJ.J		✓	✓
Shoprite Hldgs Ltd	SHPJ.J		✓+	✓
Remgro Ltd	REMJ.J		✓	✓
Aspen Pharmacare Hldgs	APNJ.J			✓
Tiger Brands Ltd	TBSJ.J			✓
ArcelorMittal SA Ltd	ACLJ.J			
Steinhoff Inter. Hldgs	SHFJ.J		✓	✓
Truworths International	TRUJ.J			✓
Mediclinic International Ltd	MDCJ.J			✓
Massmart Hldgs Ltd	MSMJ.J			✓+
Mondi Plc	MNPJ.J			
Mondi Ltd	MNDJ.J			
Imperial Hldgs Ltd	IPLJ.J			✓
Pik n Pay Stores Ltd	PIKJ.J			
Woolworths Hldgs	WHLJ.J			✓
The Foschini Grp Ltd	TFGJ.J			
Netcare Ltd	NTCJ.J			
Pretoria Port Cement	PPCJ.J			
Sappi Ltd	SAPJ.J			
Telkom SA Ltd	TKGJ.J			
Aveng Ltd	AEGJ.J			